Naval Research Laboratory

Washington, DC 20375-5320

Ocean PHILLS Data Collection and Processing: May 2000 Deployment, Lee Stocking Island, Bahamas

ROBERT A. LEATHERS
T. VALERIE DOWNES
WILLIAM A. SNYDER
JEFFREY H. BOWLES
CURTISS O. DAVIS
MARY E. KAPPUS[1]
MEGAN A. CARNEY[2]
WEI CHEN
DAN KORWAN
MARCOS J. MONTES
W. JOE RHEA

Optical Sensing Section
Radio, IR, Optical Sensors Branch
Remote Sensing Division

[1] *Current address: Observera,*
4451 Brookfield Corporate Drive, Suite 107, Chantilly, VA 20151

[2] *Current address: College of Oceanic and Atmospheric Sciences,*
Oregon State University, Corvallis, OR 97331

February 27, 2002

i

OCEAN PHILLS DATA COLLECTION AND PROCESSING: MAY 2000 DEPLOYMENT, LEE STOCKING ISLAND, BAHAMAS

INTRODUCTION

The Ocean Portable Hyperspectral Imager for Low-Light Spectroscopy (Ocean PHILLS) is a hyperspectral imager specifically designed for imaging the coastal ocean. It was deployed on an Antonov AN-2 aircraft operated by Bosch Aerospace, Inc (Huntsville, AL) in a region near Lee Stocking Island (LSI), Bahamas, in May 2000. The deployment was part of a multi-institution oceanographic study of the region sponsored by the U. S. Office of Naval Research Coastal Benthic Optical Properties program (CoBOP) [1]. The focus of CoBOP is the study of the interaction of light with the seafloor and benthic organisms in various environments. The goal is to develop techniques for using optical measurements to map and monitor water column optical properties, bathymetry, sediment types, seagrass beds, coral reefs, and other benthic species. Our objective is to use the hyperspectral PHILLS imagery along with the detailed in-situ measurements obtained during the CoBOP experiment to develop algorithms for benthic characterizations.

This document describes the LSI 2000 PHILLS data set and the manner in which it was processed to obtain remote-sensing reflectance images for use by the scientific community. This includes descriptions of spectral calibration, stray light correction, radiometric calibration, atmospheric correction, and geocorrection. Specific data file names and computer program names are provided as they existed in December 2001.

Processing of this data set proved to be very challenging due to various instrumentation difficulties that were present in 2000. The Ocean PHILLS has been in development for the past several years, and most of the problems encountered in 2000 have been corrected for subsequent deployments. The difficulties in processing the 2000 data set have been satisfactorily overcome, giving us several days of excellent remote sensing scenes over the entire LSI study site. Several research groups that participated in the CoBOP program are using this data set, and we expect this to lead to a number of journal publications (e.g., Ref [2]).

LSI 2000 DATA SET

The Ocean PHILLS, which we will hereafter simply refer to as the PHILLS, is a push-broom scanning instrument. The cross-track ground pixels are imaged with a camera lens onto a spectrometer entrance slit, and the aircraft's forward motion is used to sequentially acquire new lines of ground pixels in the along-track direction. The light passing through the entrance slit is dispersed by the spectrograph into a two-dimensional detector array to obtain a spectrum for each cross-track ground pixel. A detailed description of the PHILLS is provided in Refs. 3 and 4.

The PHILLS used at LSI (now known as PHILLS-1) measures 512 spectral channels (400 nm to 1000 nm) at 1024 spatial cross-track samples. For LSI field data and radiance calibration data, the spectral channels were binned by 4 on the CCD chip to yield 128 channels of output, each approximately 4.5 nm wide. In the laboratory, measurements taken over many frames are averaged over frames (time) to produce either a 1024 x 512 or 1024 x 128 data matrix, depending on the measurement. When flown in an aircraft, each frame covers a different cross-track line on the ground, and the result is a 1024 x 128 x (# of lines) data cube in BIL format. Generally the data is stored in sequences containing 1024 lines, which, when put together form a data "run."

The LSI 2000 data was collected using a 25 mm lens (Fujinon CF25B) at f/4 and focused at infinity, which was a setting of 5 m for this lens mounted on the PHILLS system. Data was collected at a frame rate of 28 frames/s. The aircraft flew nominally at an altitude of 2700 m and a speed of 90 knots, giving a spatial resolution on the ground of approximately 1.25 m. The camera was set in high gain mode.

All PHILLS data files recorded at LSI in 2000 are flipped with respect to sample number from the correct orientation. Therefore all files are flipped to the proper sample orientation during the calibration process. Because the camera software uses the 14 most significant bits for the data, all files are bit-shifted to make it easier to view the data.

Remote sensing data was collected near LSI on the mornings of May 16, 17, 19, 20, 21, 22, and 27, 2000 and on the afternoon of May 20, 2000. Data was also collected near Andros Island on May 24, 2000. A total of over 126 GB of data was collected during the deployment. Each data sequence contains 1024 samples, 128 spectral channels, and 1024 frames (except for the last sequence in each run, which generally contains fewer than 1024 frames) and takes 256 MB of disk space. Each pass of a flight line is given a run number and each sequence is named according to that run number and the sequence number within that run, i.e., RunXXSeqYY.bil. Corresponding to each data file are a header file (RunXXSeqYY.hdr) and a file that gives the time that each frame was collected (RunXXSeqYY.snc). For each data run there is an associated C-MIGITS file (RunXX.gps) that contains gps information along with aircraft pitch, roll, and heading. Also for each data run there is an associated dark run of at least one sequence in length that is typically taken immediately after each flight line is flown. Dark runs were taken with the lens cap on the instrument and following normal data collection procedures.

Table 1 lists all of the data collected at the LSI 2000 deployment and provides the flight line name and associated dark file for each pass of the aircraft. The raw data was initially stored on two (nearly) identical sets of 4 mm tapes. In addition, the data has been copied to digital linear tape (DLT). Table 2 shows the contents of these tapes. "Quick Look" RGB images of the data may be viewed at the Code 7212 website [5].

The flight lines for the LSI 2000 deployment are shown in Figure 1—Figure 3, superimposed on an IKONOS satellite image (Space Imaging, Thornton, CO). Figure 1 shows the AM flight lines as determined from GPS readings aboard the aircraft on the morning of May 17. Figure 2 shows the PM and closure (CL) flight lines flown in the afternoon of May 20 and for the closure experiment of May 21, respectively. The Andros flight lines of May 24 are shown in Figure 3. Andros run 08 is not shown in Figure 3 because there was no GPS information for this run; however, it lies approximately mid-way between Andros run 06 and Andros run 10. The ground coverage of the PHILLS data is illustrated in Figure 4, which shows images derived from the most commonly flow flight lines. The AM lines in Figure 4 are from May 17, which was the most cloud-free day of the deployment. Line AM4 extends beyond what is shown in this figure to a deep-water site.

Table 1: Flight lines flown and the associated data run names.

Data Run	Flight Line	Corresponding Dark Run	Data Run	Flight Line	Corresponding Dark Run
5/16/2000			**5/21/2000**		
run 02	AM 1	run 03	run 00	AM 3 (W to E)	run 01
run 04	AM 2	run 05	run 02	CL RG (W to E)	run 03
run 07	AM 3	run 08	run 04	CL CAL (W to E)	run 05
run 09	AM 4	run 10	run 06	CL RG (W to E)	run 07
run 11	AM 5	run 10	run 08	CL CAL	run 09
5/17/2000			run 10	CL RG	run 12
run 00	AM 1 (W to E)	run 01	run 11	CL RG	run 12
run 02	AM 2 (W to E)	run 03	run 13	AM 3	run 14
run 04	AM 3 (E to W)	run 05	**5/22/2000**		
run 06	AM 4 (E to W)	run 07	run 01	AM 5 (W to E)	run 03
run 08	AM 5 (E to W)	run 09	run 04	AM 1 (W to E)	run 05
run 10	AM 6 (E to W)	run 12	run 06	AM 2 (W to E)	run 07
run 11	Darby Isl.	run 12	run 08	AM 3 (W to E)	run 09
5/19/2000			run 10	AM 4 (W to E)	run 11
run 00	AM 1 (W to E)	run 01	run 12	AM 5 (W to E)	run 13
run 02	AM 2 (W to E)	run 03	run 14	AM 6 (W to E)	run 15
run 04	AM 3 (W to E)	run 05	run 16	AM 2 (W to E)	run 17
run 06	AM 4 (W to E)	run 07	**5/24/2000**		
run 08	AM 5 (W to E)	run 09	run 01	Hayburn Cay	run 02
5/20/2000 AM			run 04	Andros 6	run 05
run 01	CL RG (W to E)	run 02	run 06	Andros 5	run 07
run 03	CL RG (W to E)	run 04	run 08	Andros 4	run 09
run 05	CL CAL (W to E)	run 06	run 10	Andros 3	run 11
run 07	CL RG (W to E)	run 08	run 12	Andros 2	run 13
run 09	CL CAL (W to E)	run 10	run 14	Andros 6	run 15
run 11	CL RG (W to E)	run 12	run 16	Andros 7	run 17
run 13	CL CAL (W to E)	run 14	**5/27/2000**		
run 15	CL CAL (W to E)	run 16	run 01	AM 1	run 02
run 17	CL CAL (W to E)	run 16	run 03	AM 2	run 04
5/20/2000 PM			run 05	AM 3	run 06
run 00	PM 3	run 01	run 07	AM 4	run 08
run 02	PM 4	run 03	run 09	AM 5	run 10
run 04	CL RG	run 06	run 11	suncoaster (E to W)	run 10
run 05	CL RG new	run 06	run 12	pool	run 14
run 07	CL CAL	run 08	run 13	AM 3	run 14
run 09	CL RG new	run 10	run 15	Exuma Chain	run 17
run 11	CL CAL	run 12	run 16	continue run 15	run 17
run 13	CL RG new	run 14	run 19	AM 3	run 17

Table 2: Raw data storage locations on 4 mm tapes and DLT.

| Contents (size)	tape set 1 physical label	tape set 2 physical label	DLT label
16 May, Runs 0-11 (9.5G)	LSI_051600A	LSI 051600A(2)	LSI2000 #1
17 May, Runs 0-12 (11.8G)	LSI051700A	LSI051700A(2)	LSI2000 #1
19 May, Runs 0-15 (10.2G)	LSI051900A	LSI051900A(2)	LSI2000 #1
20 May AM, Runs 0-5, 8-10, 15-17 (11.8G)	LSI052000A	LSI 052000A(2) ver	LSI2000 #2
20 May AM, Runs 6-7, 11-14 (6.1G)	LSI052000B *	LSI052000B(2)	LSI2000 #2
20 May PM, Runs 0-14 (12.2G)	PMLSI052000Averif	PMLSI052000A(2)	LSI2000 #2
21 May, Runs 0-14 (12.7G)	LSI052100A	LSI052100A(2)	LSI2000 #3
22 May, Runs 0-9 (7.0G)	LSI052200Aver	LSI052200A(2)	LSI2000 #3
22 May, Runs 10-17 (7.0G)	LSI052200Bver	LSI052200B(2)	LSI2000 #3
24 May, Runs 6-17 (12.0G)	LSI052400A Andros	LSI052400Acy2	LSI2000 #3
24 May, Runs 0-5 (4.0G)	LSI052400B Andros	LSI052400B(2)ver	LSI2000 #3
26 May, Run 00 only (no dark) * (1.6G)	LSI052600A *	LSI052600Acopy2 *	LSI2000 #4
27 May, Runs 0-10 (10.6G)	LSI052700Aver *	LSI052700Acopy2	LSI2000 #4
27 May, Runs 11-19 (9.5G)	LSI052700Bver *	LSI052700Bcopy2	LSI2000 #4

* = no C-Migits
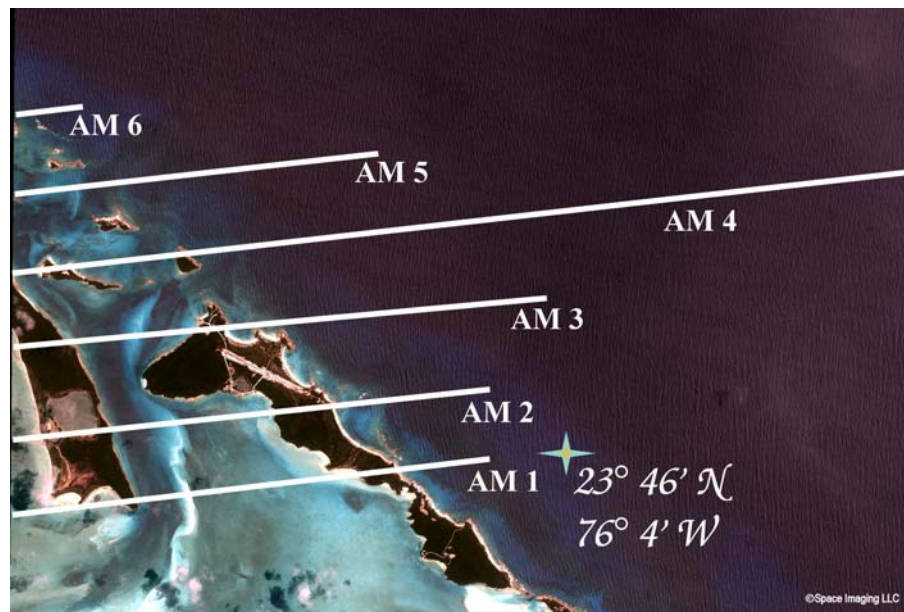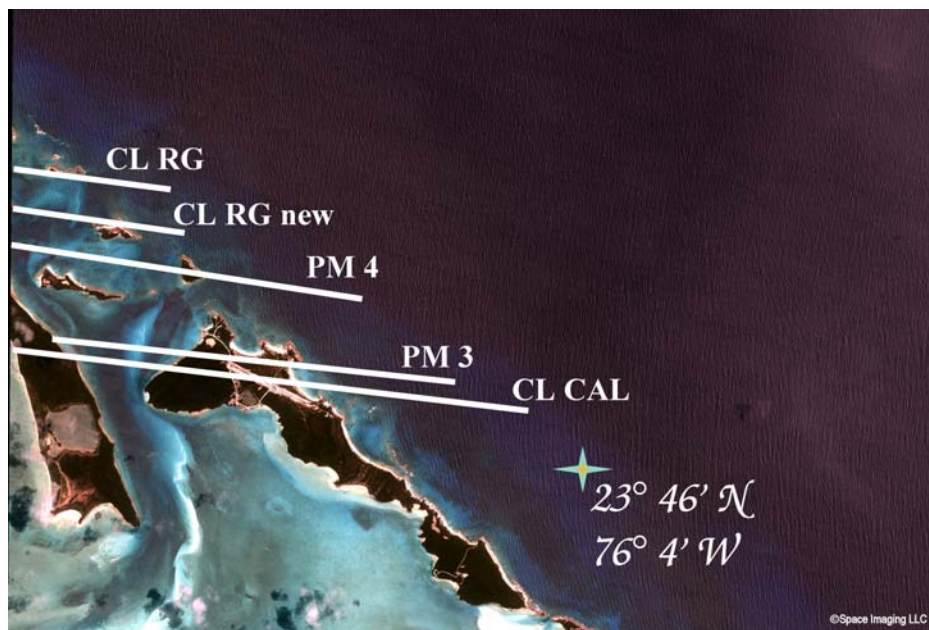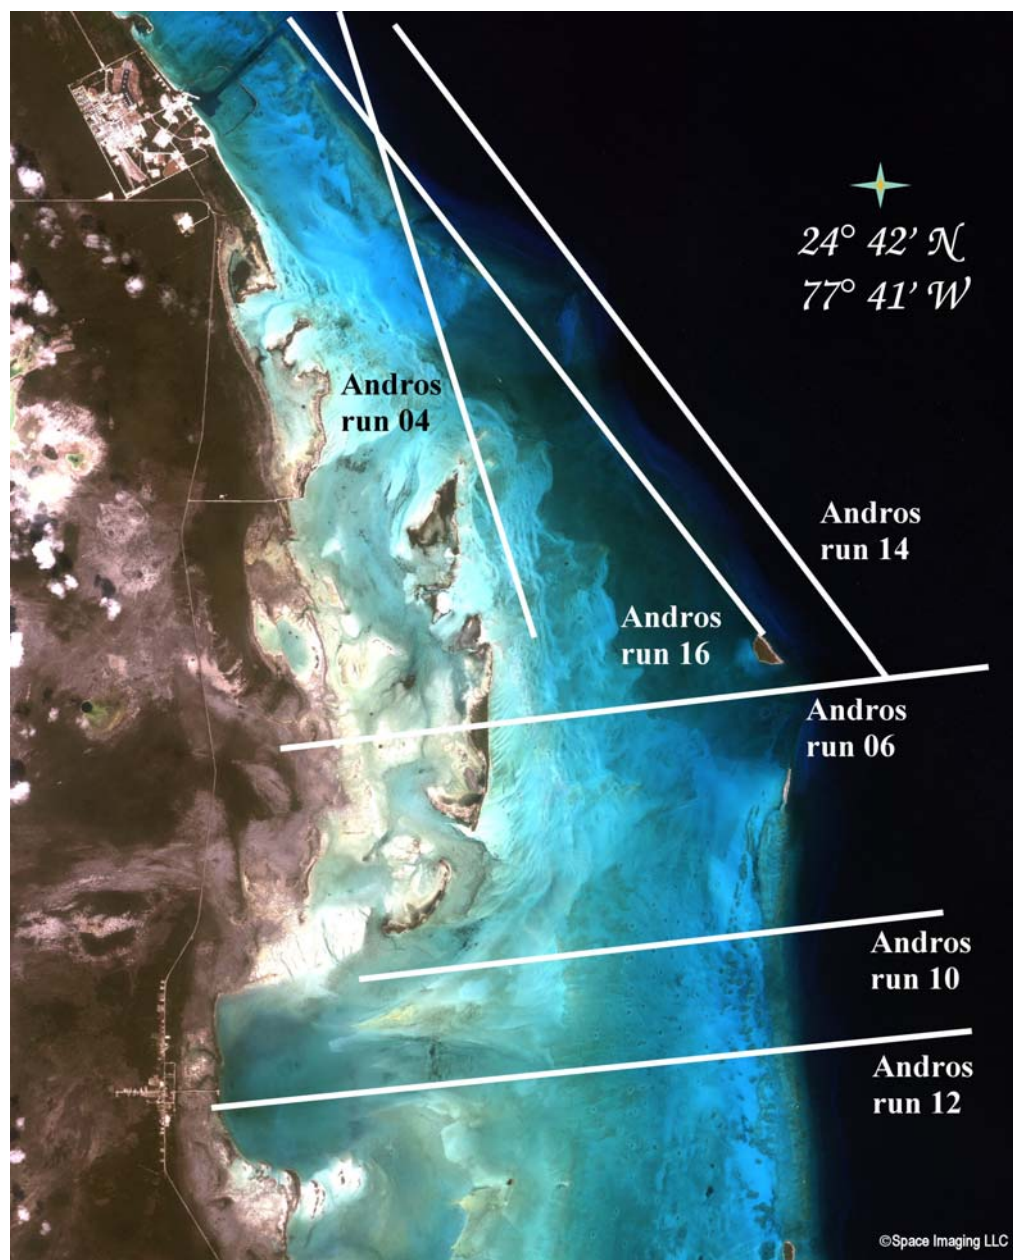
Figure 1: AM flight lines superimposed on IKONOS image of LSI.

Figure 2: PM and CL flight lines.

Figure 3: Andros flight lines.

AM 1	
AM 2	
AM 3	
AM 4	
AM 5	
AM 6	
CL CAL	
CL RG	

Figure 4: Ground coverage of the PHILLS for commonly flown flight lines.

PRE-FLIGHT SPECTRAL CALIBRATION

A spectral calibration of PHILLS-1 was performed in the laboratory in April 2000 by imaging several different gas emission lamps (oxygen, mercury, argon, helium). Data cubes were acquired measuring 512 spectral channels by imaging each gas lamp for several seconds. The cubes were averaged to provide a low noise measurement of a full spectrum at each sample. The amount of spectral smile (cross-sample variation in the spectral locations of the channels) was found to be very small, and therefore it was possible to average the gas spectra over all samples. The resulting spectrum for each lamp exhibits several emission lines (Figure 5). The exact wavelengths of known emission lines were obtained using the atom.exe program obtained from NIST. By pairing up measured emission lines with known lines (Table 3), a nearly linear relationship was derived between channel number k and wavelength of the center of each channel λ_k. The resulting fit was

$$\lambda_k \text{ (nm)} = 381.7267 + 1.2287\,k - 3.8067 \times 10^{-5}\,k^2. \qquad (1)$$

A flow chart showing the relevant file and program names for the spectral calibration is provided in Figure 6.
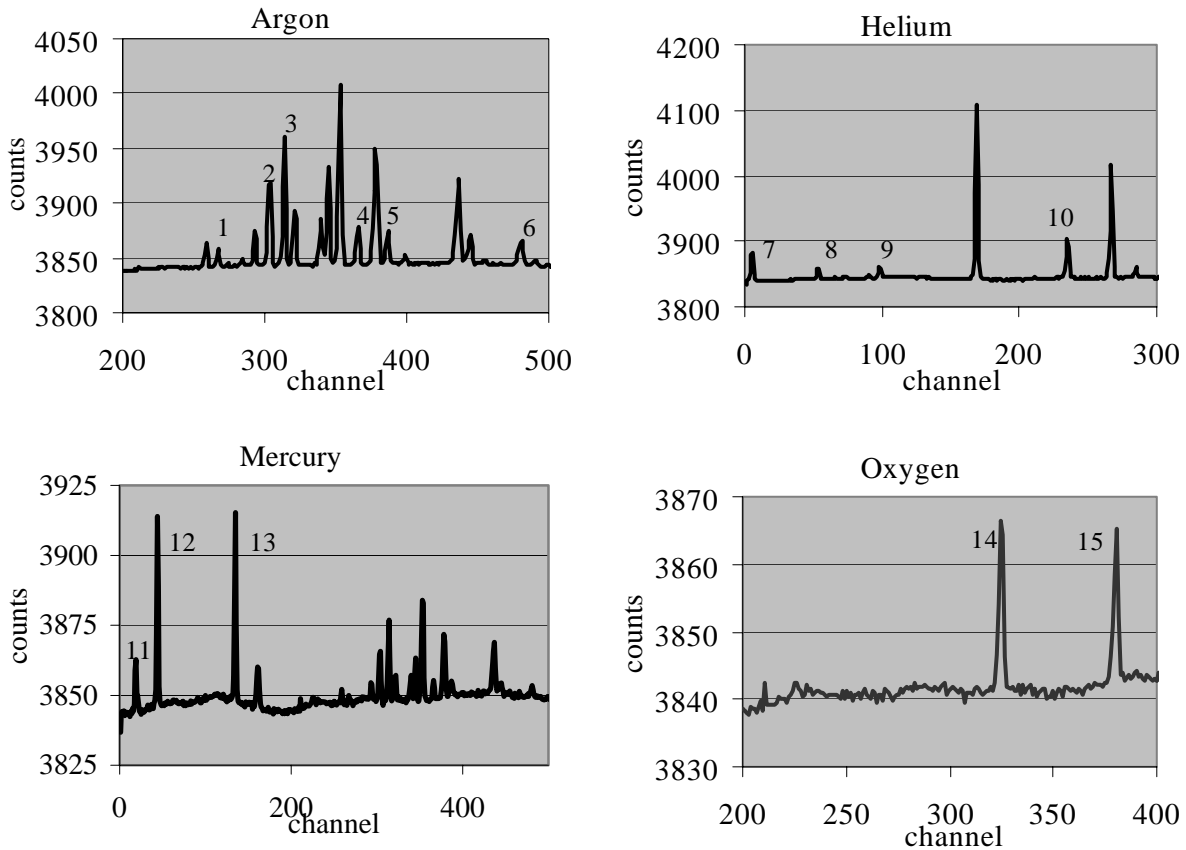

Figure 5: Spectra of gas emission lamps measured with the PHILLS. The numbered gas lines were used for spectral calibration and correspond to the line numbers in Table 3. The Hg lamp spectrum shows some Ar lines.

Table 3: Gas lines used in the April 2000 spectral calibration and their computed channel locations.

line	gas	wavelength in vacuum (nm)	channel
1	Ar	696.735	258.854
2	Ar	738.601	293.252
3	Ar	763.721	313.851
4	Ar	826.679	365.709
5	Ar	852.378	386.867
6	Ar	966.044	483.806
7	He	388.975	5.51684
8	He	447.273	53.4854
9	He	501.708	98.2567
10	He	667.999	235.189
11	Hg	404.77	18.5343
12	Hg	435.955	44.1728
13	Hg	546.226	134.916
14	O	777.631	325.319
15	O	844.868	380.766

```
┌─────────────────────┐
│  raw gas lamp data  │
│     (RunXX.bil)     │
└─────────────────────┘
           │
           ▼
╭─────────────────────╮
│ average and bit-shift│
│     (av_bs.pro)     │
╰─────────────────────╯
           │
           ▼
┌─────────────────────┐
│ sample-dependent gas│
│  emission line spectra│
│   (e.g., He_avg.bsq) │
└─────────────────────┘
           │
           ▼
╭─────────────────────╮
│  average over samples &│
│   locate & fit gas lines│
│ (spec_cal_peak_find1.pro)│
╰─────────────────────╯
           │
           ▼
┌─────────────────────┐
│ locations of selected│
│      gas lines      │
│ (spec_dat_april_A.txt)│
└─────────────────────┘
           │
           ▼
╭─────────────────────╮
│ curve fit gas line locations &│
│  interpolate using SigmaPlot │
│    (spec_april_A.jnb)   │
╰─────────────────────╯
           │
           ▼
┌─────────────────────┐
│  spectral calibration│
│ (spec_april_A512.wav,│
│  spec_april_A128.wav)│
└─────────────────────┘
```
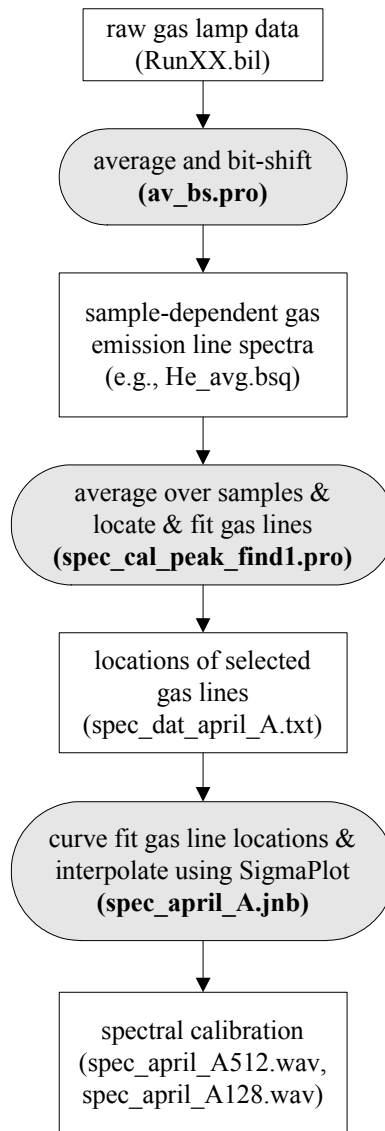
Figure 6: Spectral calibration procedure for LSI 2000.

SPECTRAL STRAY-LIGHT CORRECTION

Not all the light that enters the PHILLS is collected in the correct spectral channel. As illustrated in Figure 7, the spectral response of the PHILLS is approximately Gaussian in the channels near the wavelength being illuminated. Additionally, there is a small amount of scattered light that reaches all channels in an approximately uniform manner. This is typical of single grating monochromator-type instruments. We refer to this latter light as spectral stray light. Spectral stray light becomes particularly important whenever the intensity of the incoming light is much stronger in some parts of the spectrum than in others, in which case the small percentage of light that strays from intense parts of the spectrum causes a significant error in darker parts of the spectrum. For water spectra this typically occurs at wavelengths greater than 700 nm, while for our laboratory radiometric calibration source significant contamination takes place at both blue and red wavelengths. The Gaussian spectral response of the
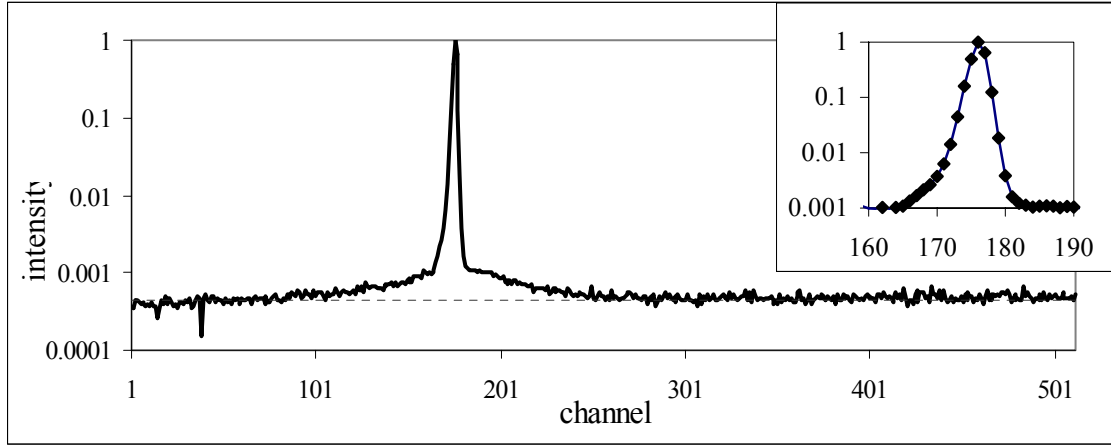
Figure 7: Spectral response of PHILLS to 594.1 nm Helium Neon laser light. The inset figure is a close up view for channels near the spectral peak. The dashed line shows the constant level of 4.3×10^{-4}.

PHILLS is accounted for during atmospheric correction; however, there are times (e.g., when deriving the zero-order correction) when this needs to be accounted for earlier in the data processing.

In what follows, a spectral correction is derived that can be used to account for the primary spectral response of the PHILLS and/or the stray light. If we at first assume that counts within each sample are conserved, then the counts measured in channel 1, c_1^m, are given by

$$
\begin{aligned}
c_1^m &= c_1^T + \left(c_2^T P_{21} + c_3^T P_{31} + \ldots \right) - \left(c_1^T P_{12} + c_1^T P_{13} + \ldots \right) \\
&= \left(c_2^T P_{21} + c_3^T P_{31} + \ldots \right) + c_1^T \left[1 - \left(P_{12} + P_{13} + \ldots \right) \right],
\end{aligned}
\tag{2}
$$

where c_1^T are the channel-1 true counts (i.e., counts that should be counted in channel 1) and P_{12} is the probability that a photon that should be counted in channel 1 instead lands in channel 2, etc. Solving for the true counts,

$$
c_1^T = \frac{\left[c_1^m - \left(c_2^T P_{21} + c_3^T P_{31} + \ldots \right) \right]}{\left[1 - \left(P_{12} + P_{13} + \ldots \right) \right]}.
\tag{3}
$$

This means that to find the true counts for a given channel we take the measured counts and subtract the counts that should not be there and then divide by the fraction of counts that should be there but aren't. Each spectral channel is described by a relationship analogous to Eq. (3). For a 3-channel system,

$$
\begin{aligned}
c_1^T \left[1 - \left(P_{12} + P_{13} \right) \right] + \left(c_2^T P_{21} + c_3^T P_{31} \right) &= c_1^m, \\
c_2^T \left[1 - \left(P_{21} + P_{23} \right) \right] + \left(c_1^T P_{12} + c_3^T P_{32} \right) &= c_2^m, \\
c_3^T \left[1 - \left(P_{31} + P_{32} \right) \right] + \left(c_1^T P_{13} + c_2^T P_{23} \right) &= c_3^m.
\end{aligned}
\tag{4}
$$

Rewriting this in matrix form,

$$
\begin{bmatrix}
1-P_{12}-P_{13} & P_{21} & P_{31} \\
P_{12} & 1-P_{21}-P_{23} & P_{32} \\
P_{13} & P_{23} & 1-P_{21}-P_{32}
\end{bmatrix}
\begin{bmatrix} c_1^T \\ c_2^T \\ c_3^T \end{bmatrix}
=
\begin{bmatrix} c_1^m \\ c_2^m \\ c_3^m \end{bmatrix},
\tag{5}
$$

the solution of which is

$$
\begin{bmatrix} c_1^T \\ c_2^T \\ c_3^T \end{bmatrix}
=
\begin{bmatrix}
1-P_{12}-P_{13} & P_{21} & P_{31} \\
P_{12} & 1-P_{21}-P_{23} & P_{32} \\
P_{13} & P_{23} & 1-P_{31}-P_{32}
\end{bmatrix}^{-1}
\begin{bmatrix} c_1^m \\ c_2^m \\ c_3^m \end{bmatrix}.
\tag{6}
$$

For an N-channel system this becomes

$$
\begin{bmatrix} c_1^T \\ c_2^T \\ \vdots \\ c_N^T \end{bmatrix}
=
\mathbf{A}
\begin{bmatrix} c_1^m \\ c_2^m \\ \vdots \\ c_N^m \end{bmatrix},
\tag{7}
$$

where the probability matrix A is given by

$$
\mathbf{A} =
\begin{bmatrix}
1-P_{12}-P_{13}\ldots-P_{1N} & P_{21} & \cdots & P_{N1} \\
P_1 & 1-P_{21}-P_{22}\ldots-P_{2N} & \cdots & P_{N2} \\
\vdots & \vdots & \ddots & \vdots \\
P_{1(N-1)} & P_{2(N-1)} & \cdots & P_{N(N-1)} \\
P_{1N} & P_{2N} & \cdots & 1-P_{N1}-P_{N2}\ldots-P_{N(N-1)}
\end{bmatrix}^{-1}.
\tag{8}
$$

The matrix **A** needs to be computed once for a given spectral spreading function. For the special case that all the probabilities are equal,

$$
\mathbf{A} =
\begin{bmatrix}
1-(N-1)P & P & P & \cdots & P \\
P & 1-(N-1)P & P & \cdots & P \\
P & P & 1-(N-1)P & \cdots & P \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
P & P & P & \cdots & 1-(N-1)P
\end{bmatrix}^{-1}.
\tag{9}
$$

To remove only the spectral stray light, we used Eq. (7) with **A** given by Eq. (9). When it was desired to correct for the spectral spreading due to the primary spectral response function as well as the stray light, we used Eq. (7) along with a more general form of **A** as described below.

The data illustrated in Figure 7 was obtained in a laboratory experiment on October 9, 2001 in which the PHILLS was used to image 543.5 nm (green) and 594.1 nm (yellow) HeNe laser light at 512 spectral channels. A similar experiment was performed with 632.8 nm (red) HeNe laser light on June 12, 2001. To determine the values for the matrix **A**, we selected a spectral response function from the dark-subtracted 632.8 nm data. We then curve-fit the stray-light function to determine the probability that a photon intended for a certain channel ends up one channel away, two channels away, etc. This fit was performed piece-wise, using a Gaussian fit for the peak, and an (exponential + linear) fit for the tail (Figure 8). Far from the peak, the data is clearly non-zero but too noisy get an accurate magnitude (note in Figure 7, however, that that the constant level for the 594.1 HeNe light appears to be 4.3×10^{-4}). Because the magnitude of the stray light far from the peak is important in certain applications, we assign it a value that was determined by trial and error (as discussed below). To populate **A** we centered the curve-fit function at channel 256, normalized so that the sum of the probabilities over the 512 channels equaled unity, and binned this by 4. Then, letting P_1 be the magnitude of the probability function one channel away from the peak and P_2 be the magnitude of the probability function two channels away from the peak, etc.,

$$
\mathbf{A} = \begin{bmatrix}
P_{11} & P_1 & P_2 & \cdots & P_{N-1} \\
P_1 & P_{22} & P_1 & \cdots & P_{N-2} \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
P_{N-1} & P_{N-2} & P_{N-3} & \cdots & P_{NN}
\end{bmatrix}^{-1} , \tag{10}
$$

where the diagonal terms are adjusted to give the desired normalization. It was decided to set all the diagonal terms to the same value, P_0, where

$$
P_0 = 1 - 2\sum_{i=1}^{63} P_i. \tag{11}
$$

In this case the peak is the same height everywhere; however, counts are no longer conserved within each sample because some counts are lost off the edge of the CCD. The underlying assumption is that the shape of the stray light probability function is the same for all wavelengths as quantified by the response of the PHILLS to 632.8 nm (HeNe laser) light.
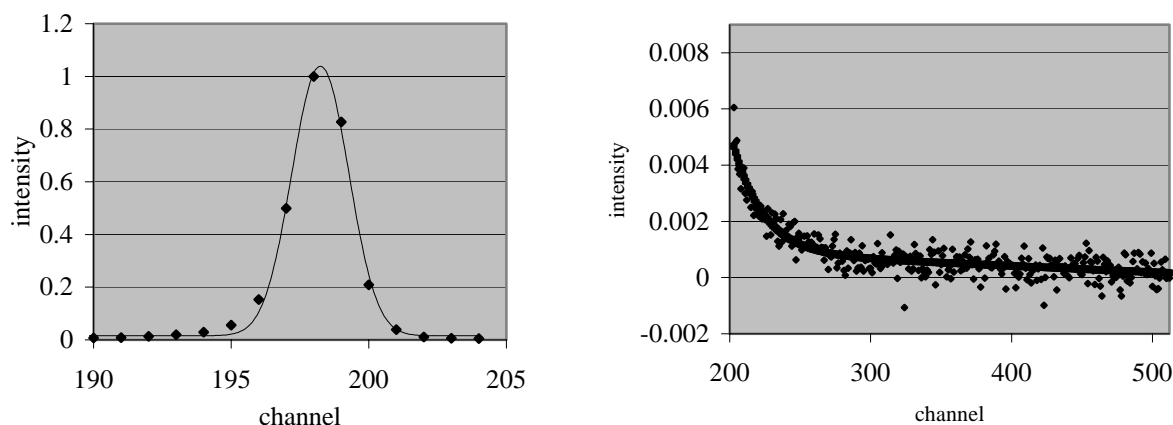
Figure 8: Curve fits of the spectral response function and stray light as quantified with normalized 633 nm HeNe data.

To determine an appropriate value for the tail of the probability function, we first defined two regions of interest in one of our data scenes that were relatively dark at red wavelengths. Using trial and error, we then applied our stray light correction to the data scene with various values for the magnitude of the tail. We chose a tail value of 4×10^{-4} (for 128 spectral channels), as it resulted in a slightly positive spectrum in the dark regions at red wavelengths. Figure 9 compares spectra from the two dark regions that are stray-light corrected with those that are uncorrected. The two regions both represent optically shallow water, one over sand covered with small marine life and the other over seagrass. The stray-light correction increases the magnitude of the peaks and decreases the spectra in the red wavelengths. It also brings together the red portions of the spectra from the two different regions, as is expected since the red light measured over water scenes is due primarily to the atmospheric contribution and skylight reflected off the sea surface, which should be the same over both regions.
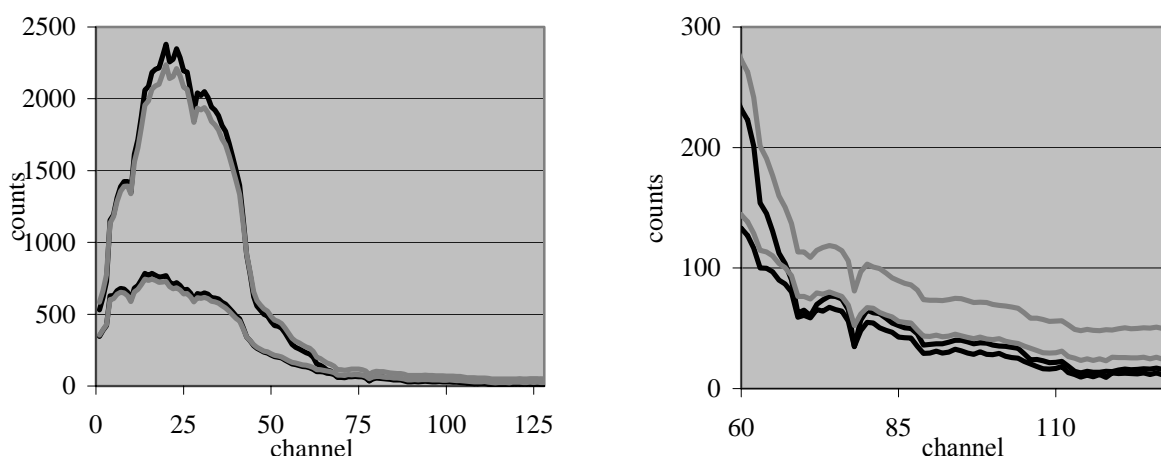

Figure 9: Comparison of stray-light corrected (black) and uncorrected (gray) spectra from red-dark scenes in May 17 (Run02Seq02). The top pair of curves is from a shallow-water scene over bio-coated sand, and the bottom pair is from shallow water over seagrass.

The flow chart in Figure 10 shows the programs and files used in the development of the spectral stray light correction.
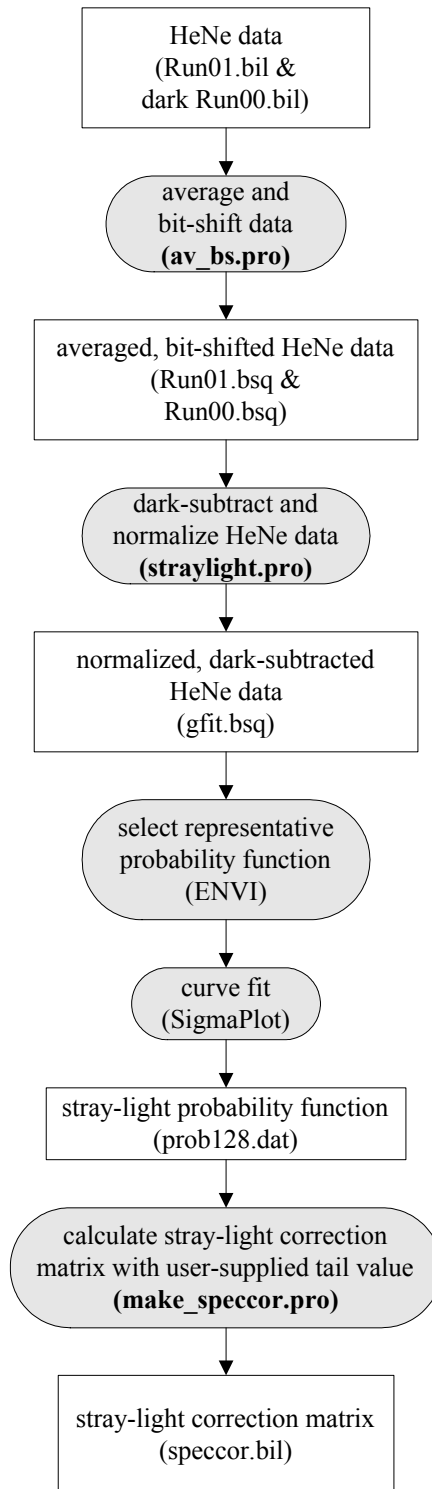

Figure 10: Procedure for generating the spectral stray-light correction from the June 12, 2001 HeNe data.

ZERO-ORDER CORRECTION

In the year 2000 deployments of the PHILLS some of the zero-order portion of the diffraction pattern was reaching the detector array and contaminating the data. This can be seen in the cross-sample profiles of the blue-wavelength channels of radiometric calibration data (Figure 11). The bumps in Figure 11 are caused by zero-order light that reflected within the camera. More specifically, this zero order light reflected from a C-mount ring surrounding the camera window and landed on the CCD array in a distinct flaring pattern, with the centers of the zero-order peaks being closer to the center of the array at high channel numbers than at low channel numbers. The measured reflectance of the C-mount ring (made of black anodized aluminum) is shown in Figure 12. Because the reflectance is very low below 680 nm, only the red wavelength portion of the zero-order light is believed to contribute significantly to the contamination. Therefore the contamination is insignificant in ocean scenes that are dark at red wavelengths. However, it is very significant in the radiometric calibration data (discussed in the next section) because the calibration is performed with a red-rich source.
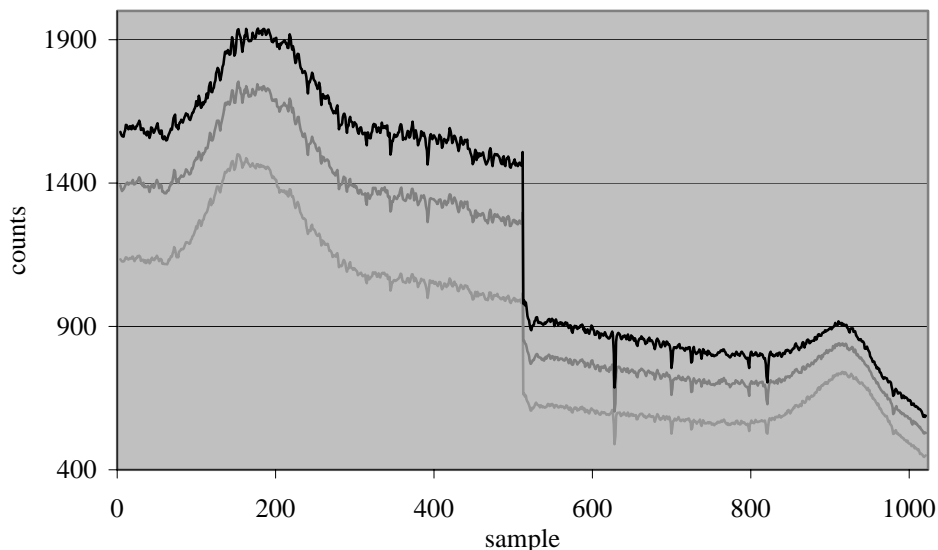

Figure 11: Cross-sample profiles of spectral channels 2, 5 and 7 (bottom to top curves) in typical year 2000 radiometric calibration data.
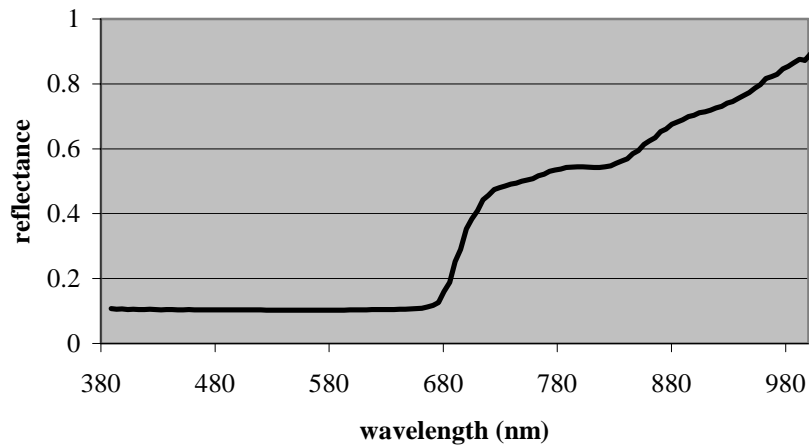
Figure 12: Reflectance of black anodized aluminum.

To help quantify the zero-order contamination, an experiment was performed in the laboratory on June 7, 2001. The PHILLS was used to image a diffuse light source (the calibration sphere) with and without a red filter placed in front of the camera lens. The red filter, made of Schott RG 630 filter glass, filters out essentially all light below 620 nm and transmits nearly all light above 700 nm. This removes essentially all of the source light at blue wavelengths but allows through most of the light that causes zero-order contamination. The light recorded by the PHILLS at wavelengths below 620 nm when the red filter was in place was therefore assumed to be either zero-order light or spectral stray light. Figure 13 shows the spectra of the filtered data at a sample near the center of the detector and at a sample near the peak of the zero order contamination. The difference between the two is due to reflected zero-order light. The magnitude of the lower curve at low channel numbers is due to a combination of stray light and zero order light.
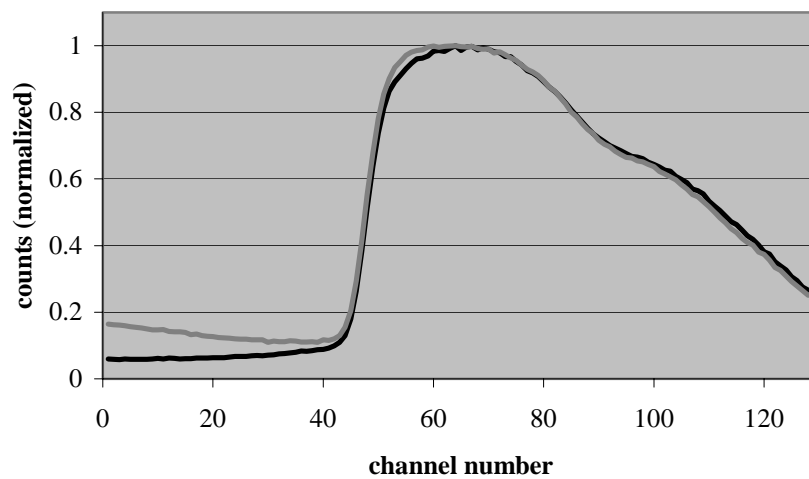

Figure 13: Red-filtered sphere data at two samples: near the center of the array (black, channel 504) and near the reflected zero-order peak (gray, channel 100).

Looking at the spectral pattern of the reflected zero-order contamination, it does not appear that any reflected zero-order light reached the center of the detector at low channel numbers. Shown in Figure 14 is the red-filtered spectrum at channel 504 (near the center of the CCD) before and after it was spectrally corrected. The corrected spectra still shows non-zero values below the filter cut-on. We considered three possible sources of these counts: 1) residual stray light due to an under-correction; 2) zero-order light reflected somewhere within the PHILLS inner casing; and 3) zero-order light reaching the CCD directly (i.e., not reflected off the inner casing). It was decided to treat this light as zero-order light reflected off the inner casing, which is predominantly black anodized aluminum. While it does not seem to fit the spatial structure of the reflected flares demonstrated in Figure 11, there is reason to believe that some zero-order light reflected off the metal-coated inactive region of the CCD array and then reflected off the inner casing back onto active areas of the CCD.
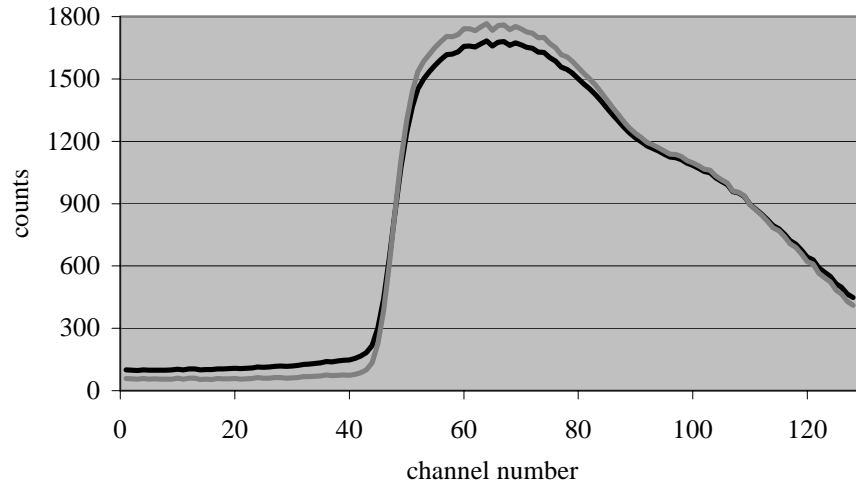

Figure 14: Uncorrected (black) and stray-light corrected (gray) curves for sample number 504 of the red-filtered sphere data.

The zero-order correction was obtained by applying the stray-light correction to the red-filtered PHILLS data and assuming that all counts remaining at channels up to the red filter cut-on (approximately channel 40) are due to zero-order light. A correction factor Z_i was obtained at each sample i and channel number j by dividing these zero-order counts $[c_0]_i$ by the sum over all channels above the red filter cut-on of the product of the (primarily first-order) counts $[c_1]_{ij}$ multiplied by the reflectance of the inner casing r_j for the wavelength corresponding to that channel.

$$Z_i = \frac{[c_0]_i}{\sum_{j=40}^{128} r_j [c_1]_{ij}} \tag{12}$$

Shown in Figure 15 is the zero order correction matrix. Cross-sample profiles of the correction matrix are shown in Figure 16. Note the spatial pattern from the reflected zero order contamination. The magnitude of the correction decreases with increasing channel number. While it appears that there exists some zero-order contamination past channel 40, it is difficult to quantify because it falls past the red filter cut-on. Although it could be estimated with extrapolation, it should be insignificant compared to the field sources being observed and was left uncorrected. The programs and data files used in computing the red-filter-based zero-order correction are given in Figure 17.
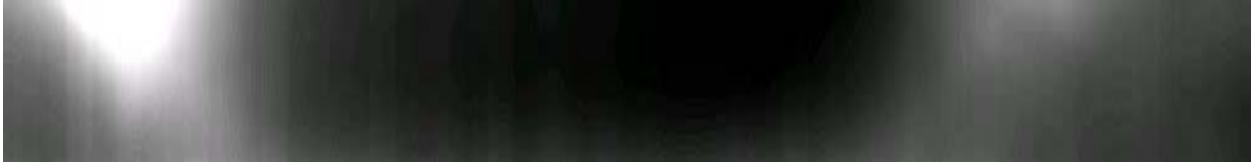
Figure 15: The zero order correction matrix versus sample (horizontal) and channel (vertical) for the first 40 (bin-by-4) channels of the detector array. Larger values are white and zero is black.
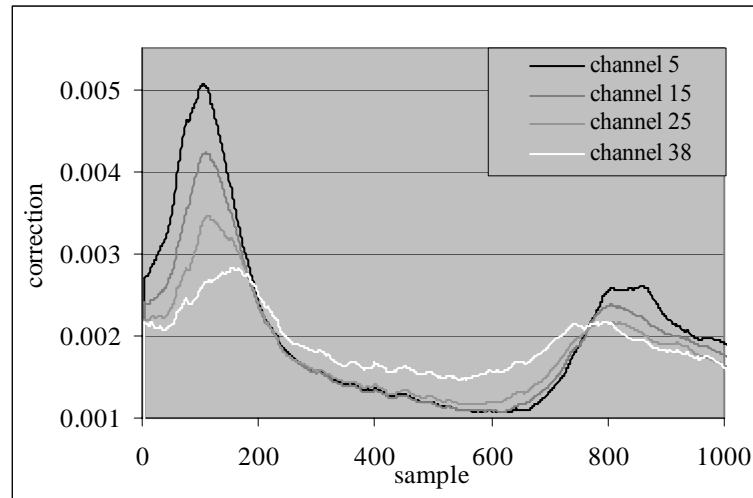

Figure 16: Example cross-sample profiles of the zero-order correction shown in Figure 15.
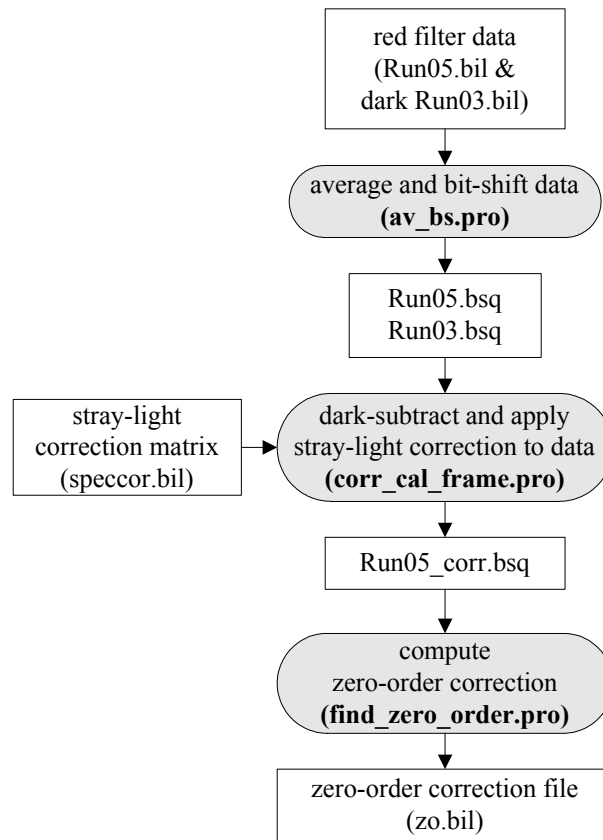
Figure 17: Flow diagram for computing the zero-order correction with red-filtered data from June 7, 2001.

The zero-order correction is applied to PHILLS data one line and one sample at a time. The zero-order counts at each sample are estimated by multiplying the correction factor for that sample by the sum over all channels in that sample of the product of the counts and the reflectance of the inner casing. The computed zero-order counts are then subtracted from the original data. Shown in Figure 18 are channels 2 and 15 of the June 2001 PHILLS image of the unfiltered sphere before and after being corrected for zero-order light. It can be seen that the zero-order bumps are successfully removed.
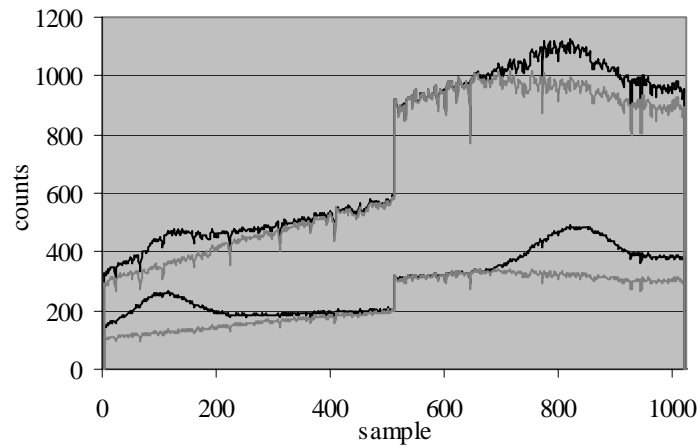
Figure 18: Channels 2 (bottom curves) and 15 (top curves) of June 2001 radiometric calibration data before (black) and after (gray) zero-order correction.

Unfortunately, the zero order correction derived from the June 2001 data does not work as well on the pre-flight radiometric calibration data collected in April 2000. As Figure 19 shows, the correction works well on the left half of the detector, but less so on the right. The change in the nature of the reflected zero-order light between 2000 and 2001 (i.e., the change in location of the right-side peak) can be seen in Figure 20. Nonetheless we believe the correction illustrated in Figure 19 is sufficient. The residual zero-order effect seen in Figure 19 is effectively eliminated during flat-fielding of the radiometric calibration coefficients (discussed in a later section).
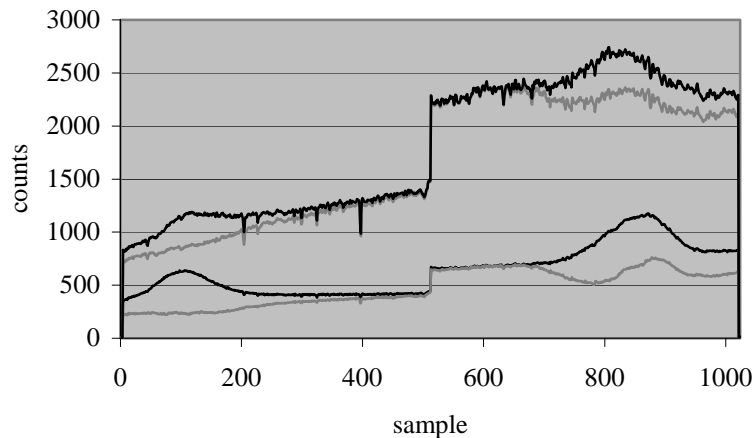

Figure 19: Channels 2 (bottom curves) and 15 (top curves) of April 2000 radiometric calibration data before (black) and after (gray) zero-order correction.
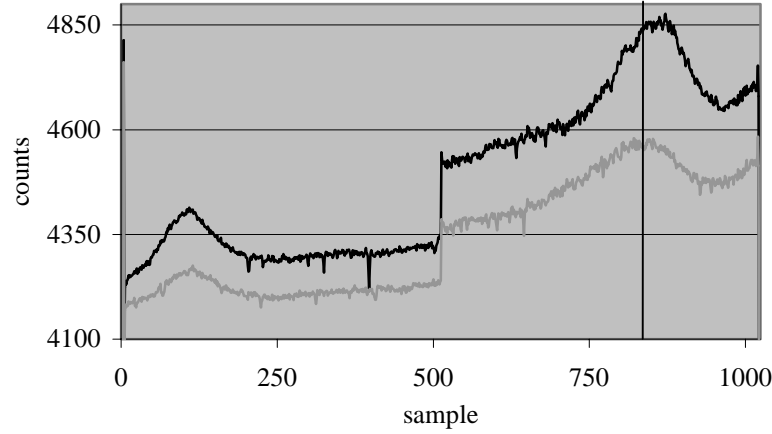
Figure 20: Cross-sample profiles of channel 3 in calibration data from April 2000 and June 2001. The vertical line shows that the right-hand side zero-order peak changed location.

Because of the shift in location of the zero-order contamination between 2000 and 2001 and because the red filter approach only works for the first 40 channels, another method of correction was also pursued. The only filtered sphere data available that was generated in the same time frame as the calibration data (May 2000) uses the blue-balancing filter (discussed in a later section). Following is the derivation of a zero-order correction based on this data.

The measured counts of filtered data for a given sample and channel (i.e., a CCD pixel), c_m^f, is a combination of first-order and zero-order light,

$$c_m^f = c_1^f + c_0^f. \tag{13}$$

Likewise for unfiltered data,

$$c_m^u = c_1^u + c_0^u. \tag{14}$$

Given the filter transmission, we can relate the first-order quantities to one another,

$$T = \frac{c_1^f}{c_1^u}. \tag{15}$$

Our zero-order correction rests on the assumption that the magnitude of the zero-order light at a given sample is proportional to the sum over all channels of the product of the housing reflectance and the zero-order counts, which in turn is proportional to the sum of the product of the housing reflectance and the first-order counts. Therefore, for channel i in a given sample,

$$\frac{\left[c_0^f\right]_i}{\left[c_0^u\right]_i} = \frac{\sum\limits_{j=1}^{128} r_j \left[c_1^f\right]_j}{\sum\limits_{j=1}^{128} r_j \left[c_1^u\right]_j} = \frac{\sum\limits_{j=1}^{128} r_j T_j \left[c_1^u\right]_j}{\sum\limits_{j=1}^{128} r_j \left[c_1^u\right]_j} \,. \tag{16}$$

Thus, for each pixel we have

$$\frac{c_0^f}{c_0^u} = f \,, \tag{17}$$

where f is approximately constant over the CCD, and can be approximated at a sample where the zero-order light is believed to be negligible (i.e., near the center of the array),

$$f = \frac{\sum\limits_{j=1}^{128} r_j T_j \left[c_1^u\right]_j}{\sum\limits_{j=1}^{128} r_j \left[c_1^u\right]_j} \,. \tag{18}$$

Note, however, that T may be sample-dependent due to geometric effects. Equations $(13) - (15)$ and (17) provide four equation and four unknowns. Solving this set of equations in terms of measured quantities, we obtain

$$c_0^u = \frac{c_m^f - c_m^u T}{f - T} \,, \tag{19}$$

$$c_0^f = \frac{f\left(c_m^f - c_m^u T\right)}{f - T} \,. \tag{20}$$

The zero-order correction can be derived either from $c_0{}^u$ or from $c_0{}^f$. The difficulty of this approach is that the spectral calibration of the PHILLS must closely match that of the spectrometer used to measure *T*. This approach shows great promise, but more work needs to be done in matching the spectral features of the filter to the PHILLS data.

PRE-FLIGHT RADIOMETRIC CALIBRATION

A pre-deployment radiometric calibration was performed in April 2000. This calibration was performed using a 40-inch Spectraflect-coated integrating sphere (Labsphere, Inc., North Sutton, NH) containing 10 halogen lamps. Measurements were made of this diffuse source with various numbers of lamps turned on, providing a range of illumination levels.

The radiance of the sphere at each lamp level is known by performing a transfer calibration from a NIST-calibrated FEL lamp. First, the known lamp irradiance is measured using an Optronics Laboratories spectroradiometer monochromator model 750-M-S with a small 6" integrating sphere (1.25" entrance port) fixed to the entrance aperture and a silicon photomultiplier tube (Si PMT) fixed to the exit slit. The monochromator is then used to measure the output from the sphere at various lamp levels (1, 2, 3, 4, 6, 8,

and 10 lamps). The radiance of the sphere is derived from the two measurements and the distance from the FEL to the input of the monochromator. The sphere intensities we used for LSI 2000 were obtained in March 2001.

The PHILLS saturated when acquiring data with four lamps illuminated, and therefore it was only possible to use unfiltered measurements at 1, 2, and 3 lamps. The sphere source is red rich and when the PHILLS saturates at red wavelengths the intensity at blue wavelengths is still below that encountered in field data. Therefore additional measurements were taken with a blue-balancing filter, which removes a portion of the red light, placed in front of the PHILLS lens. With the filter in place the intensity of the measured light is still stronger at red wavelengths than at blue wavelengths but this difference is much smaller. Measurements of the blue-filtered sphere light were made with 1, 2, 3, 4, 6, 8, and 10 lamps on. Figure 21 shows the radiance spectrum of the unfiltered sphere with 3 lamps on (the maximum possible without saturating the CCD) and that of the blue-filtered sphere with all 10 lamps on. Also shown in Figure 21 are example radiance spectra for three typical regions in LSI 2000 field data (deep water, shallow water over bright sandy bottom, and shallow water over sea grass). It can be seen that the blue-filtered spectrum is generally flatter over the spectral range of interest, which gives less sensitivity to slight errors in the spectral calibration [6]. Also, the magnitude of the filtered 10-lamp spectrum is significantly higher than the 3-lamp unfiltered spectrum below 480 nm, providing better coverage of radiance values seen in the field data at these lower wavelengths. Furthermore, the blue-filtered data remains unsaturated at more lamp levels, making it practical to use a quadratic calibration rather than a linear one. Figure 22 shows the PHILLS response to the integrating sphere with and without the blue balancing filter when three lamps are turned on.
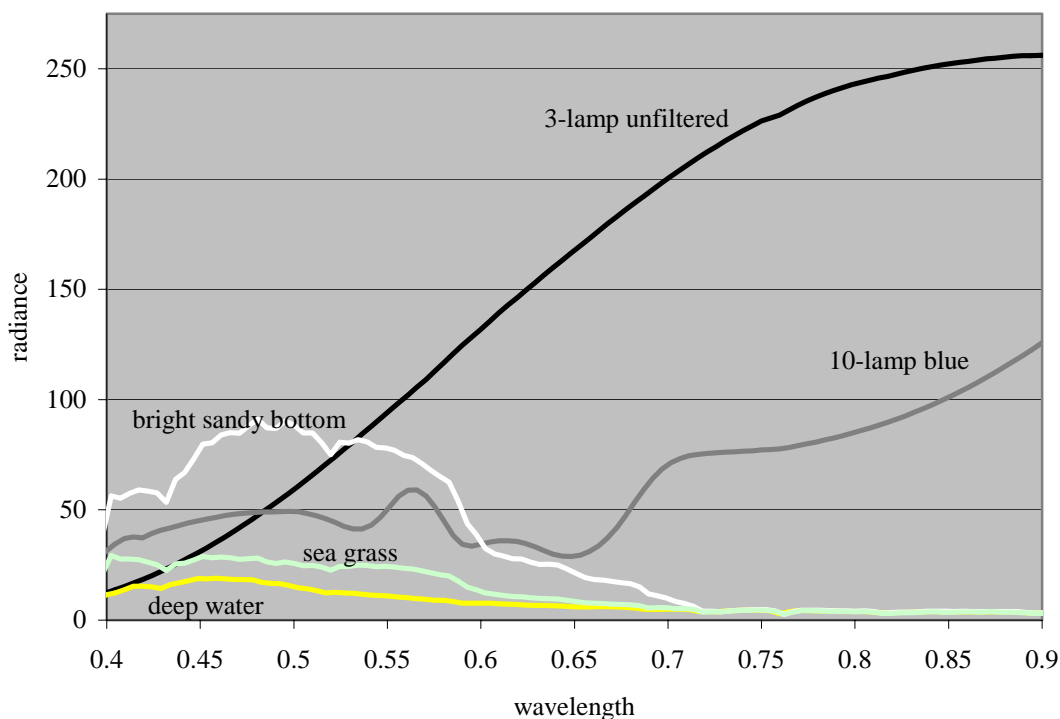

Figure 21: Radiance (W m^{-2} nm^{-1} sr^{-1}) of sphere with and without the blue filter compared with that of typical scenes from LSI 2000.
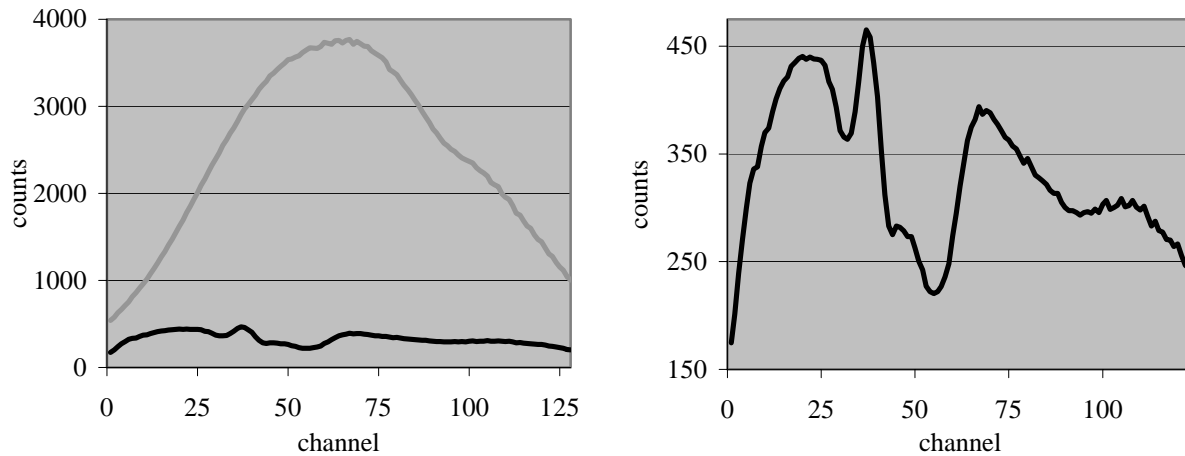
Figure 22: PHILLS response to the integrating sphere (three lamps on) with (black) and without (gray) the blue balancing filter. The right plot is a close-up of the filtered data.

The transmission of the blue balancing filter was measured in the laboratory with a Perkins Elmer Lambda 25 spectrophotometer on 17 July 2001. The transmission was measured for collimated light incident on the filter at 0, 10, and 20 degrees. The intensity of the blue-filtered sphere was computed to be that of the unfiltered sphere multiplied by the blue filter transmission interpolated to the appropriate angle (which is cross-track sample-dependent). A "pointing file" unique for each camera lens was created by measuring the angle with respect to the normal to the CCD array of the line segment between the point on the line normal to the center of the CCD array at the focal length of the lens and each sample across the CCD array. The values from the pointing file we used for the 25 mm lens (dated 12/14/00) are shown in Figure 23.
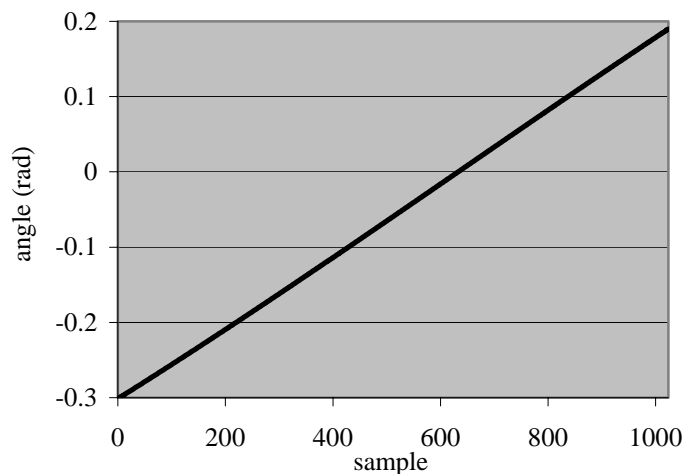

Figure 23: Pointing file for the 25 mm lens.

Figure 24 gives a summary of the radiometric calibration procedure. Table 4 lists the file names of the April 2000 calibration. For the unfiltered data it was decided to use a linear fit for each sample between the 3-lamp level and zero (lsi2000lin03.bil). The linearity of the calibration was then verified by

calibrating the one-lamp and two-lamp data and comparing the results to the correct sphere spectra. For the blue-filtered data, a quadratic fit was found using the data from 2, 4, 6, 8, and 10 lamps plus two imposed points at zero (lsi2000quad00246810.bil).
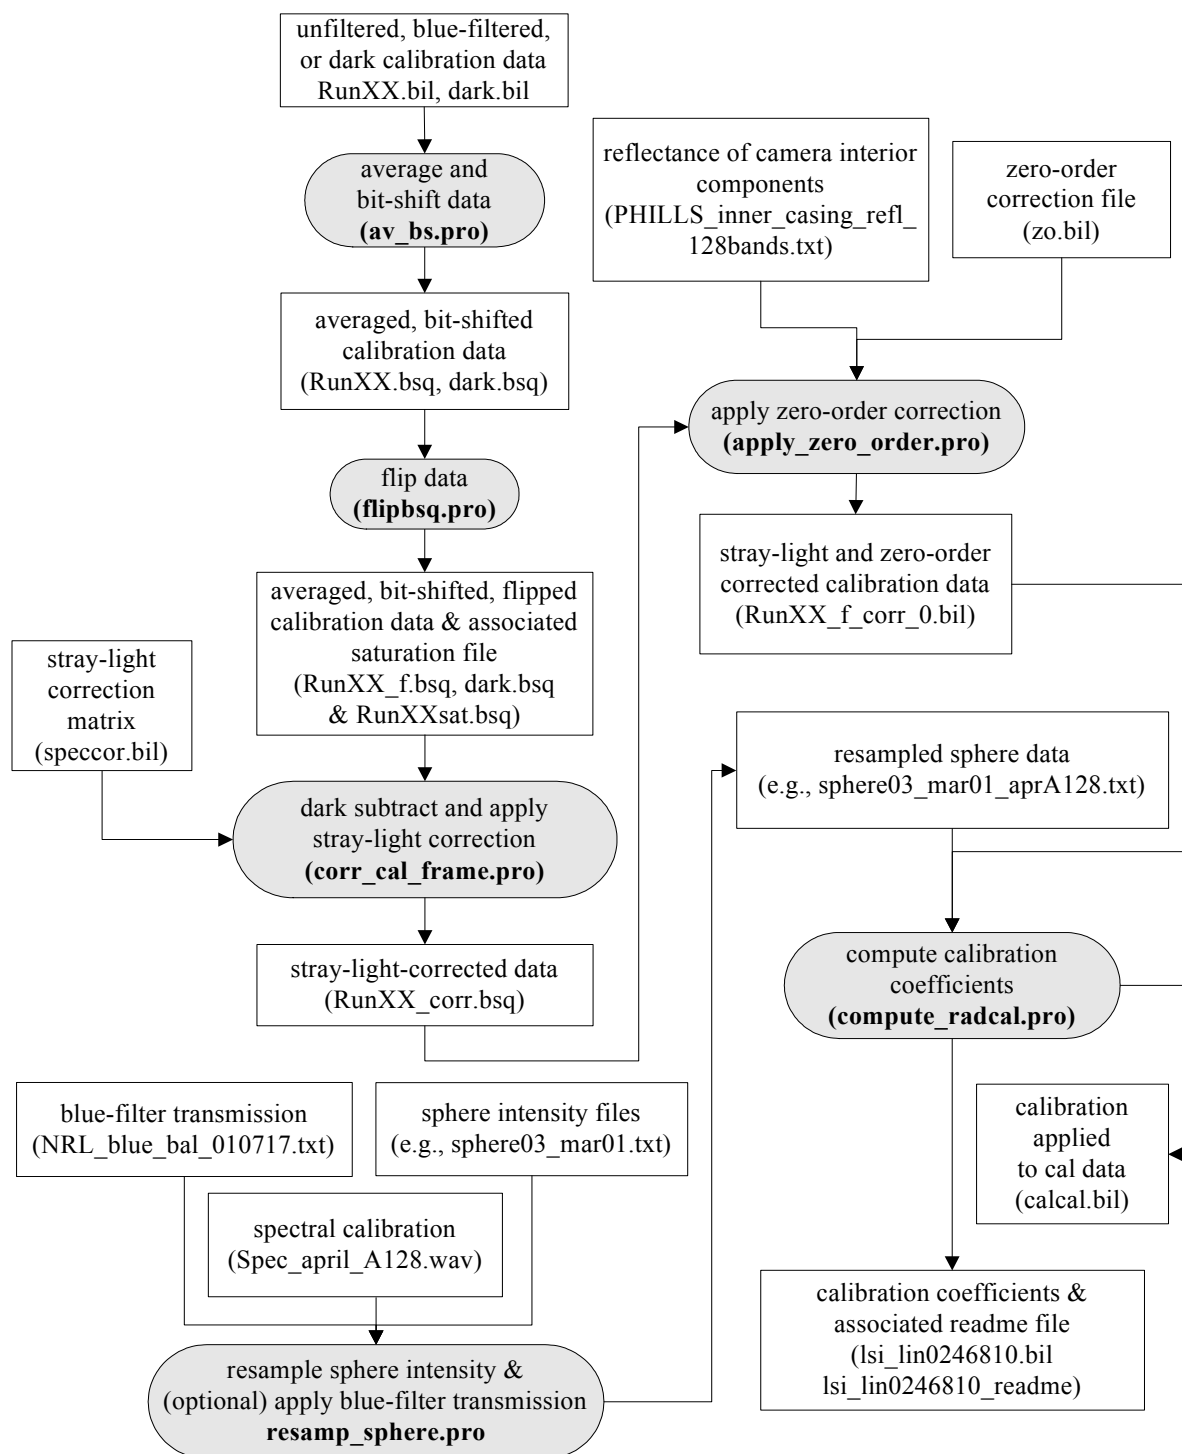

Figure 24: Radiometric calibration procedure for LSI 2000.

Table 4: File names for April 2000 radiometric calibration.

unfiltered data		
lamps	f/4.0 run	dark run
1	39	35
2	30	26
3	21	17
blue-filtered data		
lamps	f/4.0 run	dark run
1	38	35
2	29	26
3	20	17
4	15	13
6	11	13
8	7	5
10	3	1

Although the unfiltered and filtered data were found to give similar calibration coefficients, there were some small differences. The gain factor computed with the blue-filtered data was higher at high (red) wavelengths and between channels 17 and 22 than that computed with the unfiltered data. The gain factor computed with the blue-filtered data exhibits a small spectral bump near channel 42 that is likely an artifact of the blue filter itself. This is not surprising considering the shape of the filter transmittance near this channel (Figure 22).

SHIFT AND FLAT-FIELD CORRECTION OF THE RADIOMETRIC CALIBRATION

The projection of the slit in the PHILLS onto the detector array shifted by approximately 9 pixels in the sample direction between the time that the radiometric calibration was performed and the time that the field data was collected. In general, the pixel-to-pixel variation in the intensity of light reaching the detector due to nonuniformities in the slit is greater than the variations among the response of the CCD pixels. Therefore, we shifted the calibration coefficients by 9 samples before applying them to the field data. Shifting the coefficients very noticeably improves the appearance of the images. However, because some pixels in the CCD are significantly more or less sensitive than neighboring pixels, shifting the calibration coefficients introduces some vertical lines in particular channels. For example, the pixel at sample 343 and channel 38 is relatively dark and the computed calibration coefficient for this pixel is therefore relatively bright. Shifting the calibration coefficients by 9 pixels leaves pixel 343 dark in the field data and also introduces a bright pixel at sample 334.

To account for the variation in the CCD pixel-to-pixel variation, we applied a correction to the radiometric calibration coefficients through a flat-fielding approach. We selected a uniform deep-water scene (e.g., May 16 Run09Seq07) and averaged many lines together to form a set of sample-dependent spectra. We then averaged these spectra together to obtain a single spectrum that is taken to be the true value for the uniform scene. The ratio of this spectrum to each of the sample-dependent spectra gives the sample- and channel- dependent correction for the radiometric calibration. Figure 25 shows an example deep-water scene before and after flat-fielding. Figure 26 shows the flat-fielding correction obtained from the uniform scene May 16 Run09Seq07 applied to a scene from the following day.
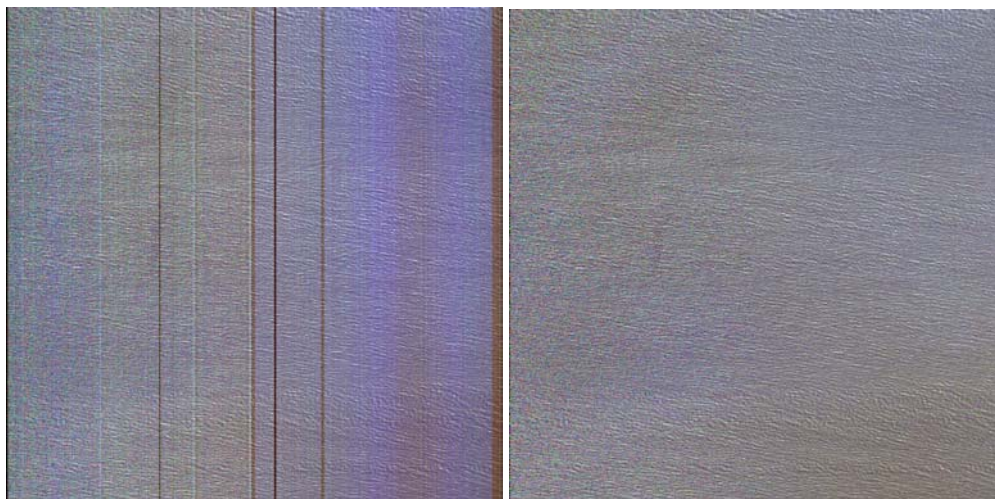
Figure 25: Deep-water scene from May 16 (Run09Seq07) before and after flat-field correction.
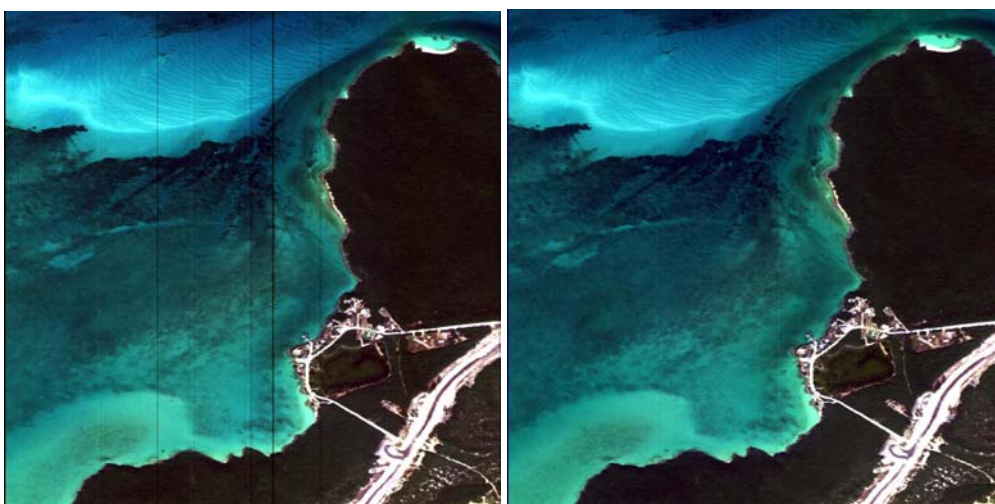

Figure 26: Sequence from May 17 (Run02Seq02) with and without flat-field correction. The correction was obtained from the scene shown in Figure 25.

The spectral response of the PHILLS is not smooth. For example, a small sawtooth pattern can be seen in the PHILLS images of the calibration sphere (e.g., Figure 22). This is believed to be an etalon effect within the CCD. It is expected that this pattern would be evident in the uncalibrated field data but that it would be removed by the radiometric calibration. However, as discussed later, we found sawtooth variation in spectra of calibrated field data. This might indicate that the radiometric calibration coefficients should also be shifted with respect to channel. This was not done, as it was decided to smooth the final remote sensing reflectance spectra instead (discussed in a later section).

APPLICATION OF CALIBRATIONS TO FIELD DATA

The steps used to calibrate field data to radiance values is outlined in Figure 27. Each field sequence needs to be flipped, stray-light corrected, and zero-order corrected. Then the radiometric calibration, flat-field correction, and spectral calibration are applied.
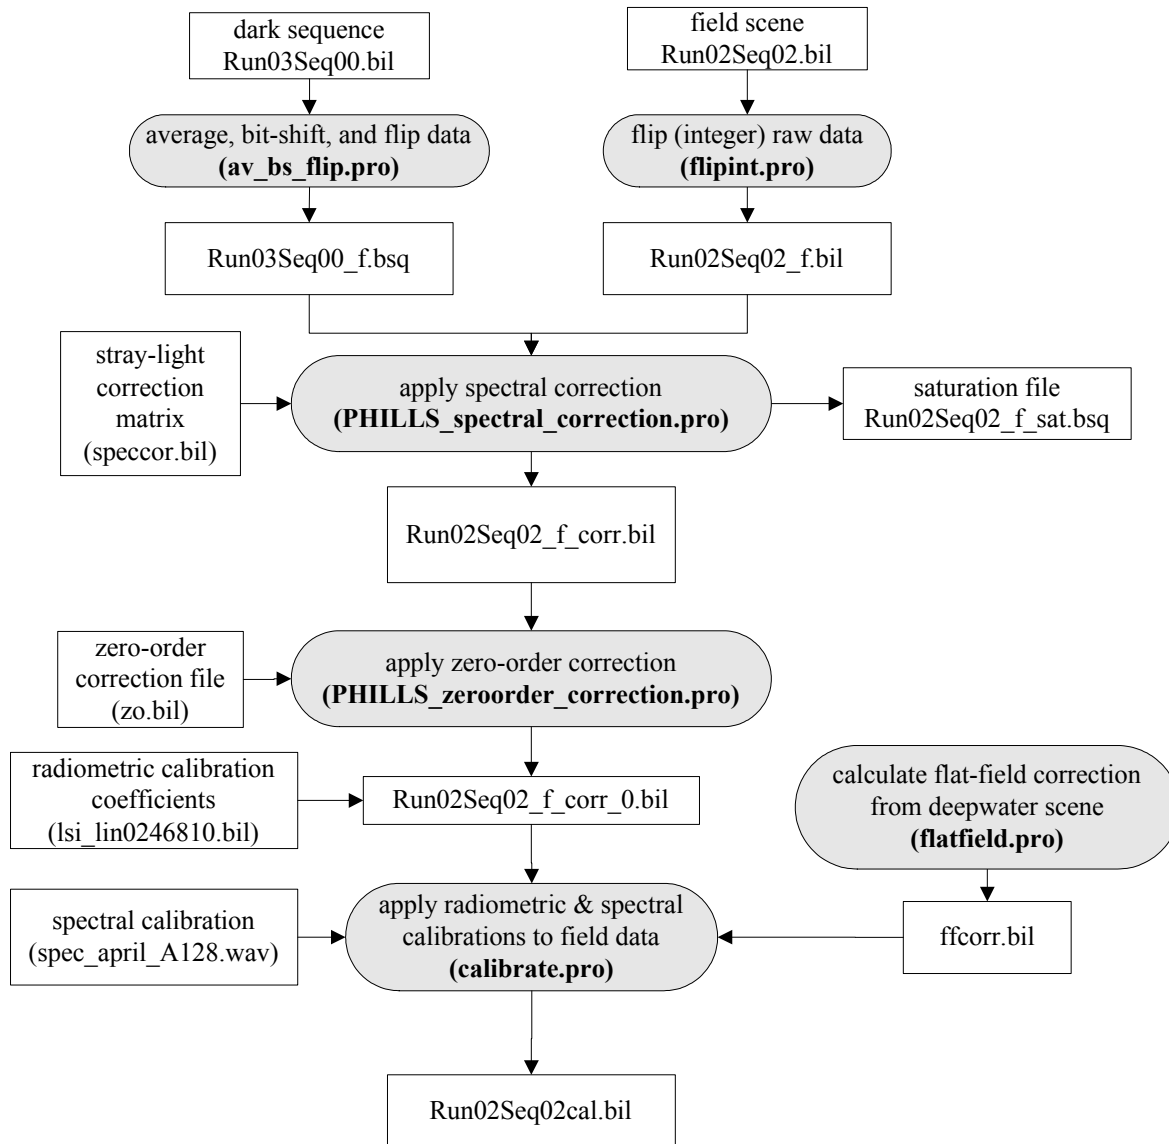

Figure 27: Calibration of the field data. The sequence May 17 Run02Seq02 is used as an example.

ATMOSPHERIC CORRECTION AND SPECTRAL ADJUSTMENTS

The radiance data was atmospherically corrected to produce remote sensing reflectance using the NRL-developed atmospheric correction algorithm Tafkaa, which is an improved version of ATREM [6]. Aerosol and gaseous scattering in Tafkaa is determined from lookup tables. There are separate tables for use over water and use over land. Tables for use over water include the wind-speed dependent specular reflection off the ocean surface due to capillary waves as determined with a Cox-Munk distribution. A description of the tables and the general algorithm may be found in Gao, et al. [7]. The application of Tafkaa is discussed in Montes, et al. [8].

The input parameters used for the atmospheric correction of a particular image are provided in the resulting header file. In particular, the user provides Tafkaa the latitude, longitude, time, aircraft altitude, relative humidity, aerosol type, aerosol optical depth (at a reference wavelength), water vapor column amount, and wind speed. Tafkaa takes into account the spectral response function of the sensor. Currently this response is modeled as a Gaussian with a specified full width half maximum.

Because Tafkaa uses different tables and algorithms for data taken over water than it does for data taken over land, it requires a land map, which we call a mask. The mask is a file that has the same number of samples and lines as the data and contains values of 100 (land) and 0 (water). To generate this mask file we first compute the NDVI (Normalized Difference Vegetation Index) for each pixel in the data set. Then by visual comparison between the PHILLS image and the NDVI file, we select the NDVI value that best separates land from water and use this cut-off to generate the mask file.

Tafkaa requires an accurate radiometric calibration and spectral calibration. The atmospheric correction is very sensitive to the spectral calibration (the channel wavelengths provided in the header file of the data file). This can be problematic in obtaining spectrally smooth remote sensing reflectance spectra. On the other hand, this can be used to help refine the spectral calibration of the field data. The channel wavelengths appear to have changed slightly between the April 2000 laboratory spectral calibration and the May 2000 field data. We adjusted the spectral calibration of the field data by shifting the wavelengths by small amounts and iteratively applying Tafkaa to see which channel wavelengths gave the smoothest remote sensing spectra. In other words, we spectrally match up the specific atmospheric absorption bands in the field data with the known location of these absorption bands used in Tafkaa. Because the Fraunhofer line at 431 nm and oxygen absorption peak at 762 nm are easiest to work with, the final spectral calibration was shifted and linearly stretched to achieve good agreement near these wavelengths. For example, channel 78 was shifted from 761.0 nm to 761.9 nm, which changed the pattern of the spectra near this channel from a dip-peak combination to a small dip only. The residual dip is a characteristic of the atmospheric removal routine over dark targets.

The steps of the atmospheric correction procedure are summarized in Figure 28.

It should be noted that there is a slight sample dependence (caused by smile and misalignment between the camera and spectrometer) in the spectral response of the PHILLS. The effect was considered minor (~ 1 mm), and it was ignored for processing the LSI 2000 data. As a result, some minor artifacts resulting from the atmospheric correction may be present near the edges of the scene.
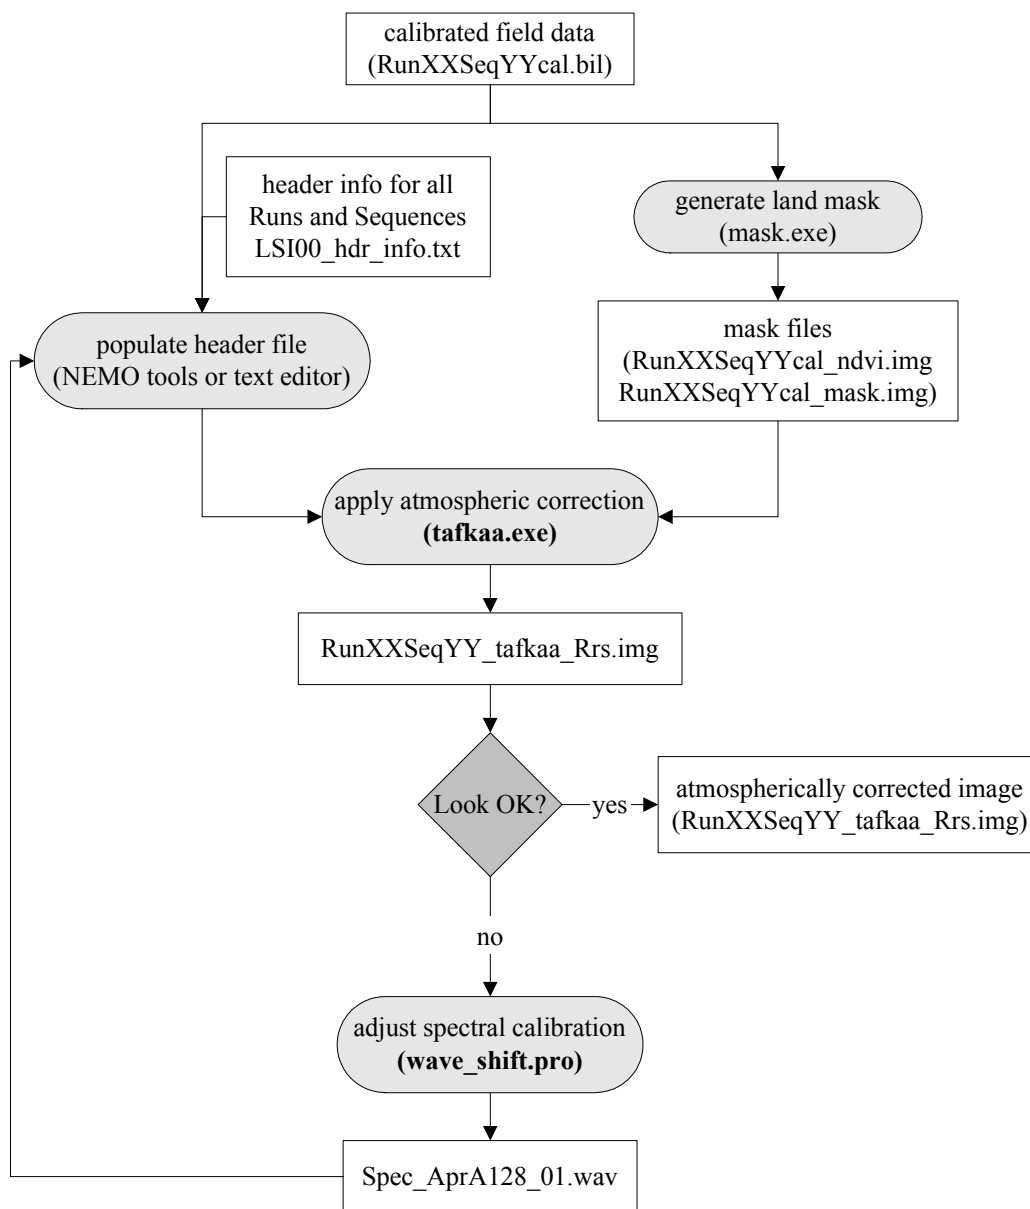
Figure 28: Atmospheric correction procedure for calibrated PHILLS data. Final spectral calibration and atmospheric water content used by Tafkaa are determined by trial and error.

SPECTRAL SMOOTHING

Binned-by-4 PHILLS spectra exhibit a slight channel-to-channel jaggedness that is not removed by the radiometric calibration. It is believed that this pattern is at least partially due to an etalon effect within the CCD. Most of the etalon effects can be removed by binning the spectral channels by 8 (rather than by 4). Alternatively, we can smooth the final remote sensing spectra. Shown in Figure 29 is an example spectrum (from May 17 Run02Seq02) before and after smoothing with IDL's averaging routine using an averaging window of 5 channels.
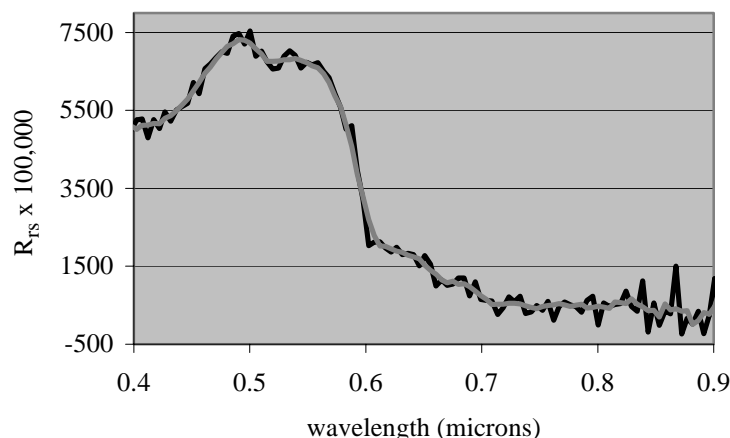
Figure 29: An example remote sensing reflectance spectrum before (black) and after (gray) smoothing.

GEOCORRECTION AND GEOREGISTRATION

During the flights the aircraft's GPS position, altitude, pitch, roll, and heading were recorded with a C-Migits GPS/INS device (Boeing Company, Chicago, IL). This information can be used to determine the position (on the ground) of each pixel in the remote sensing image to an accuracy of approximately 5 pixels. For LSI, the C-Migits and PHILLS-1 software were on different computers. It was discovered after the deployment that the clocks appear to drift relative to each other a measurable amount during the acquisition of a single scene. It has yet to be determined if this apparent drift is real or due to the difficulty in distinguishing clock drift from in-flight pitch changes.

A mapping of the angle of the incoming light for each sample on the detector array, called a pointing file, was determined in the laboratory. The pointing file used for LSI 2000 (pointing_250.txt dated 12/14/00) is plotted in Figure 23. Given the angle from the pointing file, the position on the ground of each sample is computed from the aircraft position, altitude, pitch, roll, and heading.

Using a set of cross-track ground control points (GCP's), we determined the fixed roll, pitch, and heading offsets of the C-Migits device in the aircraft. The GCP's were acquired both from hand-held GPS information taken on the island and by matching up landscape features to pixels on a geolocated IKONOS satellite image (June 27, 2000) for the region. Once the C-Migits offsets were determined, the time offset was computed for several different lines within each sequence by matching along-track GCP's. Generally, the time offset varied linearly with respect to line number and therefore could be easily interpolated.

A flow chart of the geocorrection process is provided in Figure 30. The C-Migits data is used to generate an Input Geometry (IGM) file that provides the latitude and longitude for each pixel in the scene. Using ENVI, we use the IGM files to create Geographic Lookup Tables (GLT) files that provide the information necessary to warp the image to remove distortion due to aircraft motion. Finally, the GLT is used with the PHILLS image to produce the final product.
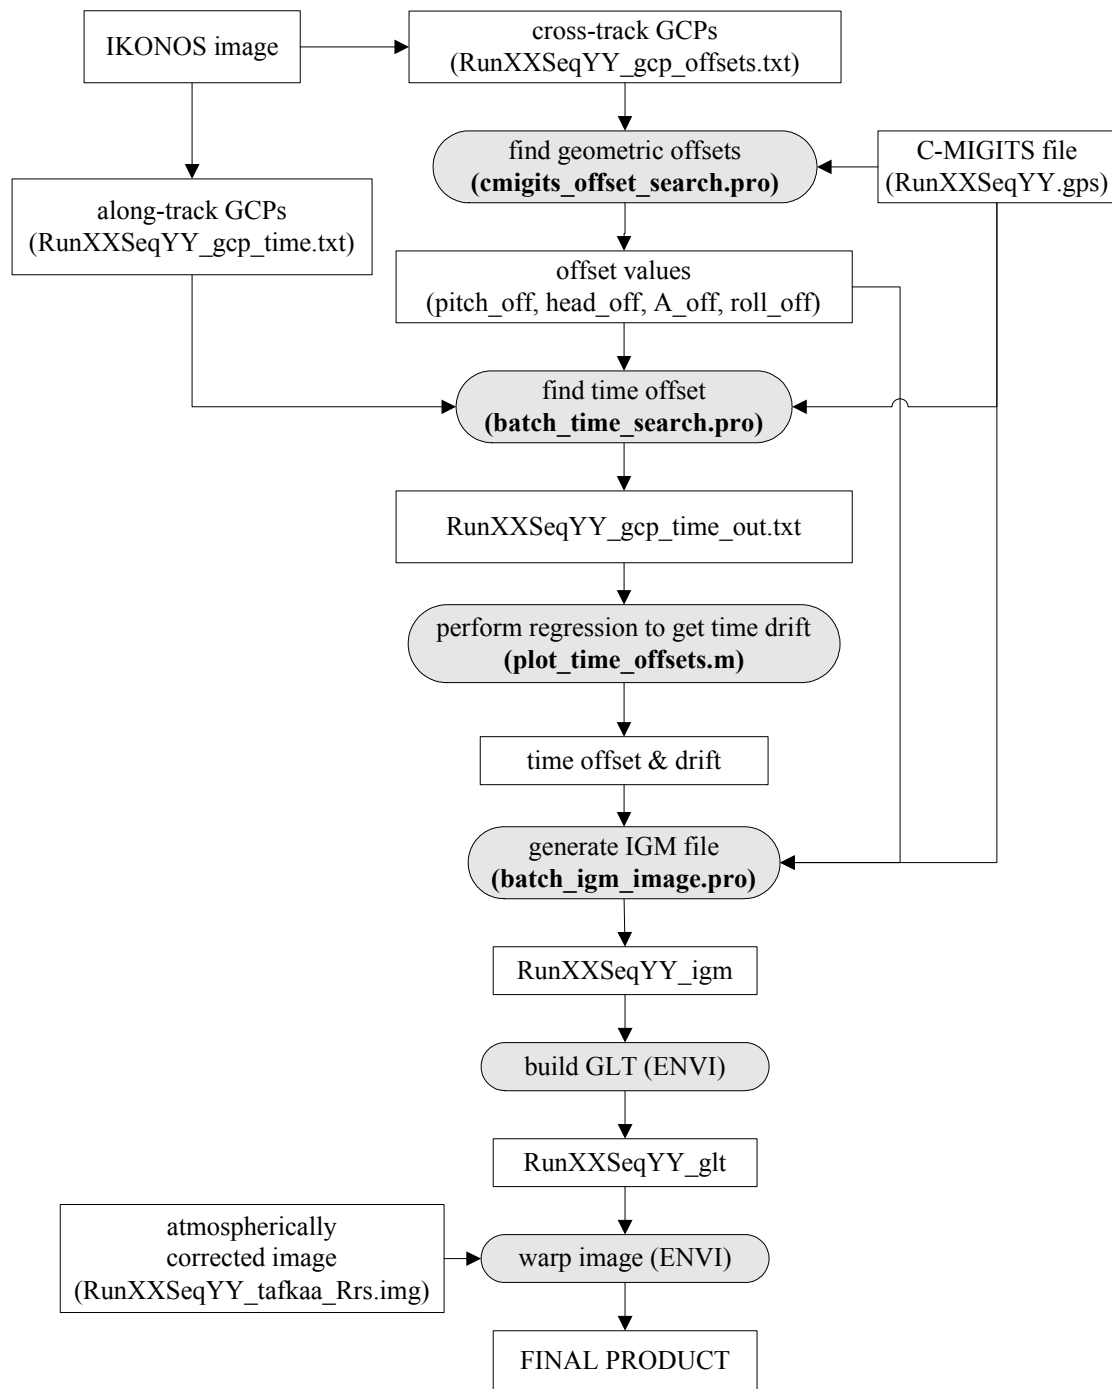
Figure 30: Example data flow for geocorrection of LSI 2000 PHILLS data.

An example geocorrected image of Adderly Cut is shown in Figure 31 along with the corresponding section of the IKONOS image. The shape of the borders of the PHILLS image shows the amount of warping that was applied to correct this image.
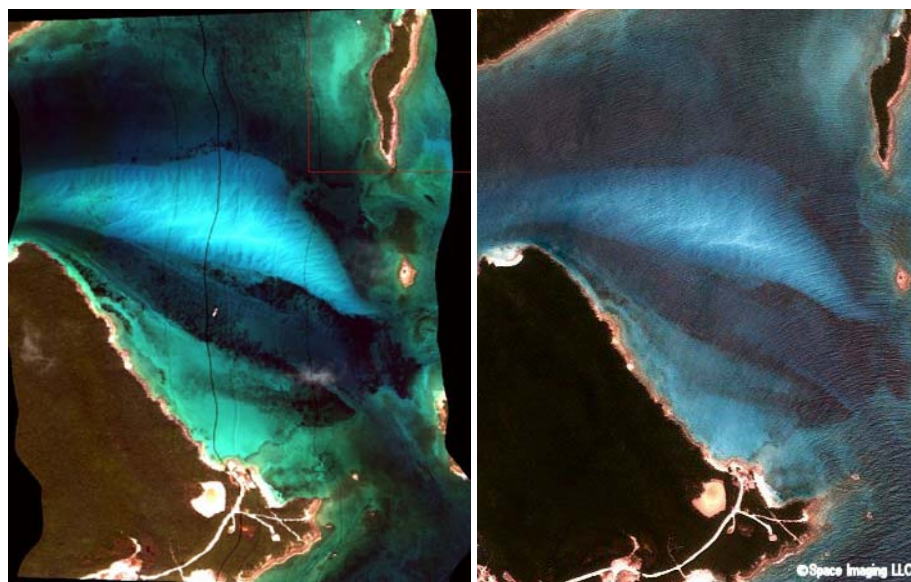
Figure 31: Geocorrected PHILLS image of Adderly Cut (left) compared with an IKONOS image of the same region (right).

REFERENCES

1. Office of Naval Research (ONR) Coastal Benthic Optical Properties (CoBOP) program, http://www.psicorp.com/cobop/cobop.html.

2. H. M. Dierssen, R. C. Zimmerman, R. A. Leathers, T. V. Downes, and C. O. Davis, "Remote sensing of seagrass and bathymetry using high resolution airborne imagery," *Limnology & Oceanography* (in review).

3. C. Davis, M. Kappus, J. Bowles, J. Fisher, J Antoniades, and M Carney, "Calibration, characterization and first results with the Ocean PHILLS Hyperspectral Imager," Proc. SPIE **3753**: 160-168, 1999.

4. C. O. Davis, J. Bowles, R. A. Leathers, D. Korwan, T. V. Downes, W. A. Snyder, W. J. Rhea, W. Chen, J. Fisher, P. Bissett, and R. A. Reisse, "Ocean PHILLS hyperspectral imager: design, characterization, and calibration," *Optics Express* **10**: 210-221, 2002.

5. NRL Optical Sensing Section, IR/Radio/Optical Sensing Branch, Remote Sensing Division. "Quick Look images from the CoBOP 2000 experiment," http://rsd-www.nrl.navy.mil/7212/lsi/LSI00QL.htm.

6. J. Bowles, M. Kappus, J. Antoniades, M. Baumback, M Czarnaski, C. O. Davis and J. Grossmann, "Calibration of inexpensive pushbroom imaging spectrometers," *Metrologica* **35**: 657-661, 1999.

7. B.-C. Gao, M. J. Montes, Z. Ahmad, and C. O. Davis, "Atmospheric correction algorithm for hyperspectral remote sensing of ocean color from space", *Applied Optics* **39**: 887-896, 2000.

8. M. J. Montes, B.-C. Gao, and C. O. Davis, "A new algorithm for atmospheric correction of hyperspectral remote sensing data," in *Geo-Spatial Image and Data Exploitation II*, W. E. Roper, ed., Proc. SPIE **4383**: 23-30, 2001.